



US011373088B2

(12) **United States Patent**  
**Bleiweiss et al.**

(10) **Patent No.: US 11,373,088 B2**  
(45) **Date of Patent: Jun. 28, 2022**

(54) **MACHINE LEARNING ACCELERATOR  
MECHANISM**

**G06N 3/08** (2006.01)  
**G06F 9/00** (2006.01)  
**G06T 1/20** (2006.01)

(71) Applicant: **Intel Corporation**, Santa Clara, CA  
(US)

(52) **U.S. Cl.**  
CPC ..... **G06N 3/063** (2013.01); **G06F 7/78**  
(2013.01); **G06F 9/00** (2013.01); **G06N 3/084**  
(2013.01); **G06N 20/00** (2019.01); **G06F**  
**2207/4824** (2013.01); **G06T 1/20** (2013.01)

(72) Inventors: **Amit Bleiweiss**, Yad Binyamin (IL);  
**Anavai Ramesh**, Chandler, AZ (US);  
**Asit Mishra**, Hillsboro, OR (US);  
**Deborah Marr**, Portland, OR (US);  
**Jeffrey Cook**, Portland, OR (US);  
**Srinivas Sridharan**, Bangalore (IN);  
**Eriko Nurvitadhi**, Hillsboro, OR (US);  
**Elmoustapha Ould-Ahmed-Vall**,  
Chandler, AZ (US); **Dheevatsa**  
**Mudigere**, Bangalore (IN);  
**Mohammad Ashraf Bhuiyan**,  
Beaverton, OR (US); **Md Faijul Amin**,  
Chandler, AZ (US); **Wei Wang**, Santa  
Clara, CA (US); **Dhawal Srivastava**,  
Phoenix, AZ (US); **Niharika**  
**Maheshwari**, Santa Clara, CA (US)

(58) **Field of Classification Search**  
CPC ..... G06N 3/063; G06N 20/00; G06N 3/084;  
G06N 3/08; G06N 3/0454; G06N 3/04;  
G06N 3/0445; G06N 7/005; G06N 5/04;  
G06N 3/02; G06N 20/20; G06N 3/0472;  
G06N 3/088; G06F 7/78; G06F 9/00;  
G06F 2207/4824; G06T 1/20; G06T  
2207/20081; G06T 2207/20084; G06K  
9/6267

See application file for complete search history.

(73) Assignee: **INTEL CORPORATION**, Santa Clara,  
CA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

9,646,243 B1 \* 5/2017 Gokmen ..... G06N 3/0454  
10,452,971 B2 \* 10/2019 Chung ..... G06N 3/063  
10,452,995 B2 \* 10/2019 Burger ..... G06N 20/00  
2006/0202038 A1 \* 9/2006 Wang ..... G06K 7/10811  
235/462.24

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 1205 days.

(Continued)

*Primary Examiner* — Abderrahim Merouan

(74) *Attorney, Agent, or Firm* — Jaffery Watson;  
Mendonsa & Hamilton LLP

(21) Appl. No.: **15/859,504**

(22) Filed: **Dec. 30, 2017**

(65) **Prior Publication Data**

US 2019/0205737 A1 Jul. 4, 2019


(51) **Int. Cl.**  
**G06N 3/063** (2006.01)  
**G06N 20/00** (2019.01)  
**G06F 7/78** (2006.01)

(57) **ABSTRACT**

An apparatus to facilitate acceleration of machine learning  
operations is disclosed. The apparatus comprises at least one  
processor to perform operations to implement a neural  
network and accelerator logic to perform communicatively  
coupled to the processor to perform compute operations for  
the neural network.

**25 Claims, 50 Drawing Sheets**

1700



(56)                      **References Cited**

U.S. PATENT DOCUMENTS

2006/0283952 A1 \* 12/2006 Wang ..... G06K 7/10  
235/462.01  
2016/0379109 A1 \* 12/2016 Chung ..... G06F 15/7803  
706/26  
2017/0124451 A1 \* 5/2017 Barham ..... G06N 3/08  
2017/0262995 A1 \* 9/2017 Li ..... G06N 3/0454  
2017/0343481 A1 \* 11/2017 Jahanshahi ..... G21D 3/001  
2017/0346792 A1 \* 11/2017 Nataros ..... G06F 13/28  
2018/0060958 A1 \* 3/2018 Gershon ..... G06Q 40/04  
2018/0144242 A1 \* 5/2018 Simard ..... G06N 3/0454  
2018/0189638 A1 \* 7/2018 Nurvitadhi ..... G06N 3/0445  
2018/0191629 A1 \* 7/2018 Biederman ..... H04L 47/6225  
2018/0191632 A1 \* 7/2018 Biederman ..... H04L 45/20  
2019/0156212 A1 \* 5/2019 Bottaro ..... G06N 5/022  
2019/0286973 A1 \* 9/2019 Kowuri ..... G06F 8/451  
2021/0125070 A1 \* 4/2021 Wang ..... G06N 3/063

\* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4

500




FIG. 5

EXECUTION LOGIC  
600




FIG. 6A




FIG. 6B

GRAPHICS PROCESSOR INSTRUCTION FORMATS




FIG. 7




FIG. 8

FIG. 9A

GRAPHICS PROCESSOR COMMAND FORMAT


900



FIG. 9B

GRAPHICS PROCESSOR COMMAND SEQUENCE

910






FIG. 10

IP CORE DEVELOPMENT - 1100




FIG. 11A




FIG. 11B




FIG. 12




FIG. 13A




FIG. 13B




FIG. 14A




FIG. 14B

1500




FIG. 15




FIG. 16A




FIG. 16B




FIG. 17




FIG. 18




FIG. 19




FIG. 20




FIG. 21A




FIG. 21B




FIG. 22




FIG. 23




FIG. 24





FIG. 25



Forward propagation:  $(M \times K) = (K \times N)$

Backward propagation:  $(M \times K)^T = (M \times N)$

Weight update:  $(M \times N) = (K \times N)^T$

*FIG. 26*








FIG. 27



*FIG. 28B*



*FIG. 28C*



*FIG. 28D*




FIG. 28E




FIG. 28F




FIG. 28G




FIG. 28H




FIG. 28I




FIG. 29




FIG. 30




FIG. 31




FIG. 32




FIG. 33A




FIG. 33B





FIG. 33C






FIG. 34A




FIG. 34B




FIG. 35A




FIG. 35B




FIG. 35C




FIG. 35D




FIG. 35E




FIG. 35F




FIG. 35G




FIG. 36

## 1

**MACHINE LEARNING ACCELERATOR  
MECHANISM**

## FIELD

Embodiments relate generally to data processing and more particularly to data processing via a general-purpose graphics processing unit.

## BACKGROUND OF THE DESCRIPTION

Current parallel graphics data processing includes systems and methods developed to perform specific operations on graphics data such as, for example, linear interpolation, tessellation, rasterization, texture mapping, depth testing, etc. Traditionally, graphics processors used fixed function computational units to process graphics data; however, more recently, portions of graphics processors have been made programmable, enabling such processors to support a wider variety of operations for processing vertex and fragment data.

To further increase performance, graphics processors typically implement processing techniques such as pipelining that attempt to process, in parallel, as much graphics data as possible throughout the different parts of the graphics pipeline. Parallel graphics processors with single instruction, multiple thread (SIMT) architectures are designed to maximize the amount of parallel processing in the graphics pipeline. In an SIMT architecture, groups of parallel threads attempt to execute program instructions synchronously together as often as possible to increase processing efficiency. A general overview of software and hardware for SIMT architectures can be found in Shane Cook, CUDA Programming Chapter 3, pages 37-51 (2013).

## BRIEF DESCRIPTION OF THE DRAWINGS

So that the manner in which the above recited features of the present embodiments can be understood in detail, a more particular description of the embodiments, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments and are therefore not to be considered limiting of its scope.

FIG. 1 is a block diagram of a processing system, according to an embodiment.

FIG. 2 is a block diagram of an embodiment of a processor having one or more processor cores, an integrated memory controller, and an integrated graphics processor.

FIG. 3 is a block diagram of a graphics processor, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a plurality of processing cores.

FIG. 4 is a block diagram of a graphics processing engine of a graphics processor in accordance with some embodiments.

FIG. 5 is a block diagram of hardware logic of a graphics processor core according to some embodiments.

FIG. 6A-6B illustrate thread execution logic including an array of processing elements employed in a graphics processor core according to some embodiments.

FIG. 7 is a block diagram illustrating a graphics processor instruction formats according to some embodiments.

FIG. 8 is a block diagram of another embodiment of a graphics processor.

## 2

FIG. 9A is a block diagram illustrating a graphics processor command format according to an embodiment.

FIG. 9B is a block diagram illustrating a graphics processor command sequence according to an embodiment.

FIG. 10 illustrates exemplary graphics software architecture for a data processing system according to some embodiments.

FIG. 11A is a block diagram illustrating an IP core development system that may be used to manufacture an integrated circuit to perform operations according to an embodiment.

FIG. 11B illustrates a cross-section side view of an integrated circuit package assembly according to some embodiments.

FIG. 12 is a block diagram illustrating an exemplary system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment.

FIGS. 13A-13B are block diagrams illustrating exemplary graphics processors for use within an System on Chip (SoC), according to embodiments described herein.

FIGS. 14A-14B illustrate additional exemplary graphics processor logic according to embodiments described herein.

FIG. 15 illustrates a machine learning software stack, according to an embodiment.

FIGS. 16A-16B illustrate layers of exemplary deep neural networks.

FIG. 17 illustrates an exemplary recurrent neural network.

FIG. 18 illustrates training and deployment of a deep neural network.

FIG. 19 is a block diagram illustrating distributed learning.

FIG. 20 illustrates a computing device employing a machine learning acceleration mechanism, according to an embodiment.

FIG. 21A illustrates an exemplary differentiable neural computer (DNC).

FIG. 21B illustrates one embodiment of a DNC flow.

FIG. 22 is a flow diagram illustrating one embodiment of a scheduling process implemented at a machine learning acceleration mechanism.

FIG. 23 illustrates one embodiment of a shortest path search implemented at a machine learning acceleration mechanism.

FIG. 24 illustrates one embodiment of a map implemented at a machine learning acceleration mechanism.

FIG. 25 illustrates one embodiment of neural network layers.

FIG. 26 illustrates one embodiment of matrices.

FIG. 27 illustrates one embodiment of a machine learning accelerator.

FIGS. 28A-28F illustrate embodiments of a machine learning accelerator.

FIG. 28G illustrates one embodiment of a processing element.

FIG. 28H illustrates one embodiment of a scratchpad.

FIG. 28I illustrates one embodiment of a functional unit.

FIG. 29 is a flow diagram illustrating a conventional batch flow process.

FIG. 30 is a flow diagram illustrating one embodiment of a batch flow process.

FIG. 31 is a flow diagram illustrating one embodiment of a convolutional-normalization process.

FIG. 32 is a block diagram illustrating a computer system configured to implement one or more aspects of the embodiments described herein.

FIG. 33A-33D illustrate parallel processor components, according to an embodiment.

## 3

FIGS. 34A-34B are block diagrams of graphics multiprocessors, according to embodiments.

FIGS. 35A-35G illustrate an exemplary architecture in which a plurality of GPUs are communicatively coupled to a plurality of multi-core processors.

FIG. 36 illustrates a graphics processing pipeline, according to an embodiment.

## DETAILED DESCRIPTION

In embodiments, various mechanisms for accelerating machine learning operations are described.

In the following description, numerous specific details are set forth to provide a more thorough understanding. However, it will be apparent to one of skill in the art that the embodiments described herein may be practiced without one or more of these specific details. In other instances, well-known features have not been described to avoid obscuring the details of the present embodiments.

## System Overview

FIG. 1 is a block diagram of a processing system 100, according to an embodiment. In various embodiments, the system 100 includes one or more processors 102 and one or more graphics processors 108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 102 or processor cores 107. In one embodiment, the system 100 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

In one embodiment, the system 100 can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In some embodiments, the system 100 is a mobile phone, smart phone, tablet computing device or mobile Internet device. The processing system 100 can also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart eyewear device, augmented reality device, or virtual reality device. In some embodiments, the processing system 100 is a television or set top box device having one or more processors 102 and a graphical interface generated by one or more graphics processors 108.

In some embodiments, the one or more processors 102 each include one or more processor cores 107 to process instructions which, when executed, perform operations for system and user software. In some embodiments, each of the one or more processor cores 107 is configured to process a specific instruction set 109. In some embodiments, instruction set 109 may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores 107 may each process a different instruction set 109, which may include instructions to facilitate the emulation of other instruction sets. Processor core 107 may also include other processing devices, such as a Digital Signal Processor (DSP).

In some embodiments, the processor 102 includes cache memory 104. Depending on the architecture, the processor 102 can have a single internal cache or multiple levels of internal cache. In some embodiments, the cache memory is shared among various components of the processor 102. In some embodiments, the processor 102 also uses an external cache (e.g., a Level-3 (L3) cache or Last Level Cache (LLC)) (not shown), which may be shared among processor cores 107 using known cache coherency techniques. A

## 4

register file 106 is additionally included in processor 102 which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). Some registers may be general-purpose registers, while other registers may be specific to the design of the processor 102.

In some embodiments, one or more processor(s) 102 are coupled with one or more interface bus(es) 110 to transmit communication signals such as address, data, or control signals between processor 102 and other components in the system 100. The interface bus 110, in one embodiment, can be a processor bus, such as a version of the Direct Media Interface (DMI) bus. However, processor busses are not limited to the DMI bus, and may include one or more Peripheral Component Interconnect buses (e.g., PCI, PCI Express), memory busses, or other types of interface busses. In one embodiment the processor(s) 102 include an integrated memory controller 116 and a platform controller hub 130. The memory controller 116 facilitates communication between a memory device and other components of the system 100, while the platform controller hub (PCH) 130 provides connections to I/O devices via a local I/O bus.

The memory device 120 can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In one embodiment the memory device 120 can operate as system memory for the system 100, to store data 122 and instructions 121 for use when the one or more processors 102 executes an application or process. Memory controller 116 also couples with an optional external graphics processor 112, which may communicate with the one or more graphics processors 108 in processors 102 to perform graphics and media operations. In some embodiments a display device 111 can connect to the processor(s) 102. The display device 111 can be one or more of an internal display device, as in a mobile electronic device or a laptop device or an external display device attached via a display interface (e.g., DisplayPort, etc.). In one embodiment the display device 111 can be a head mounted display (HMD) such as a stereoscopic display device for use in virtual reality (VR) applications or augmented reality (AR) applications.

In some embodiments the platform controller hub 130 enables peripherals to connect to memory device 120 and processor 102 via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller 146, a network controller 134, a firmware interface 128, a wireless transceiver 126, touch sensors 125, a data storage device 124 (e.g., hard disk drive, flash memory, etc.). The data storage device 124 can connect via a storage interface (e.g., SATA) or via a peripheral bus, such as a Peripheral Component Interconnect bus (e.g., PCI, PCI Express). The touch sensors 125 can include touch screen sensors, pressure sensors, or fingerprint sensors. The wireless transceiver 126 can be a Wi-Fi transceiver, a Bluetooth transceiver, or a mobile network transceiver such as a 3G, 4G, or Long Term Evolution (LTE) transceiver. The firmware interface 128 enables communication with system firmware, and can be, for example, a unified extensible firmware interface (UEFI). The network controller 134 can enable a network connection to a wired network. In some embodiments, a high-performance network controller (not shown) couples with the interface bus 110. The audio controller 146, in one embodiment, is a multi-channel high definition audio controller. In one embodiment the system 100 includes an optional legacy

## 5

I/O controller **140** for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. The platform controller hub **130** can also connect to one or more Universal Serial Bus (USB) controllers **142** connect input devices, such as keyboard and mouse **143** combinations, a camera **144**, or other USB input devices.

It will be appreciated that the system **100** shown is exemplary and not limiting, as other types of data processing systems that are differently configured may also be used. For example, an instance of the memory controller **116** and platform controller hub **130** may be integrated into a discreet external graphics processor, such as the external graphics processor **112**. In one embodiment the platform controller hub **130** and/or memory controller **160** may be external to the one or more processor(s) **102**. For example, the system **100** can include an external memory controller **116** and platform controller hub **130**, which may be configured as a memory controller hub and peripheral controller hub within a system chipset that is in communication with the processor(s) **102**.

FIG. 2 is a block diagram of an embodiment of a processor **200** having one or more processor cores **202A-202N**, an integrated memory controller **214**, and an integrated graphics processor **208**. Those elements of FIG. 2 having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. Processor **200** can include additional cores up to and including additional core **202N** represented by the dashed lined boxes. Each of processor cores **202A-202N** includes one or more internal cache units **204A-204N**. In some embodiments each processor core also has access to one or more shared cached units **206**.

The internal cache units **204A-204N** and shared cache units **206** represent a cache memory hierarchy within the processor **200**. The cache memory hierarchy may include at least one level of instruction and data cache within each processor core and one or more levels of shared mid-level cache, such as a Level 2 (L2), Level 3 (L3), Level 4 (L4), or other levels of cache, where the highest level of cache before external memory is classified as the LLC. In some embodiments, cache coherency logic maintains coherency between the various cache units **206** and **204A-204N**.

In some embodiments, processor **200** may also include a set of one or more bus controller units **216** and a system agent core **210**. The one or more bus controller units **216** manage a set of peripheral buses, such as one or more PCI or PCI express busses. System agent core **210** provides management functionality for the various processor components. In some embodiments, system agent core **210** includes one or more integrated memory controllers **214** to manage access to various external memory devices (not shown).

In some embodiments, one or more of the processor cores **202A-202N** include support for simultaneous multi-threading. In such embodiment, the system agent core **210** includes components for coordinating and operating cores **202A-202N** during multi-threaded processing. System agent core **210** may additionally include a power control unit (PCU), which includes logic and components to regulate the power state of processor cores **202A-202N** and graphics processor **208**.

In some embodiments, processor **200** additionally includes graphics processor **208** to execute graphics processing operations. In some embodiments, the graphics processor **208** couples with the set of shared cache units **206**, and the system agent core **210**, including the one or more

## 6

integrated memory controllers **214**. In some embodiments, the system agent core **210** also includes a display controller **211** to drive graphics processor output to one or more coupled displays. In some embodiments, display controller **211** may also be a separate module coupled with the graphics processor via at least one interconnect, or may be integrated within the graphics processor **208**.

In some embodiments, a ring based interconnect unit **212** is used to couple the internal components of the processor **200**. However, an alternative interconnect unit may be used, such as a point-to-point interconnect, a switched interconnect, or other techniques, including techniques well known in the art. In some embodiments, graphics processor **208** couples with the ring interconnect **212** via an I/O link **213**.

The exemplary I/O link **213** represents at least one of multiple varieties of I/O interconnects, including an on package I/O interconnect which facilitates communication between various processor components and a high-performance embedded memory module **218**, such as an eDRAM module. In some embodiments, each of the processor cores **202A-202N** and graphics processor **208** use embedded memory modules **218** as a shared Last Level Cache.

In some embodiments, processor cores **202A-202N** are homogenous cores executing the same instruction set architecture. In another embodiment, processor cores **202A-202N** are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores **202A-202N** execute a first instruction set, while at least one of the other cores executes a subset of the first instruction set or a different instruction set. In one embodiment processor cores **202A-202N** are heterogeneous in terms of microarchitecture, where one or more cores having a relatively higher power consumption couple with one or more power cores having a lower power consumption. Additionally, processor **200** can be implemented on one or more chips or as an SoC integrated circuit having the illustrated components, in addition to other components.

FIG. 3 is a block diagram of a graphics processor **300**, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a plurality of processing cores. In some embodiments, the graphics processor communicates via a memory mapped I/O interface to registers on the graphics processor and with commands placed into the processor memory. In some embodiments, graphics processor **300** includes a memory interface **314** to access memory. Memory interface **314** can be an interface to local memory, one or more internal caches, one or more shared external caches, and/or to system memory.

In some embodiments, graphics processor **300** also includes a display controller **302** to drive display output data to a display device **320**. Display controller **302** includes hardware for one or more overlay planes for the display and composition of multiple layers of video or user interface elements. The display device **320** can be an internal or external display device. In one embodiment the display device **320** is a head mounted display device, such as a virtual reality (VR) display device or an augmented reality (AR) display device. In some embodiments, graphics processor **300** includes a video codec engine **306** to encode, decode, or transcode media to, from, or between one or more media encoding formats, including, but not limited to Moving Picture Experts Group (MPEG) formats such as MPEG-2, Advanced Video Coding (AVC) formats such as H.264/MPEG-4 AVC, as well as the Society of Motion Picture & Television Engineers (SMPTE) 421M/VC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

In some embodiments, graphics processor **300** includes a block image transfer (BLIT) engine **304** to perform two-dimensional (2D) rasterizer operations including, for example, bit-boundary block transfers. However, in one embodiment, 2D graphics operations are performed using one or more components of graphics processing engine (GPE) **310**. In some embodiments, GPE **310** is a compute engine for performing graphics operations, including three-dimensional (3D) graphics operations and media operations.

In some embodiments, GPE **310** includes a 3D pipeline **312** for performing 3D operations, such as rendering three-dimensional images and scenes using processing functions that act upon 3D primitive shapes (e.g., rectangle, triangle, etc.). The 3D pipeline **312** includes programmable and fixed function elements that perform various tasks within the element and/or spawn execution threads to a 3D/Media sub-system **315**. While 3D pipeline **312** can be used to perform media operations, an embodiment of GPE **310** also includes a media pipeline **316** that is specifically used to perform media operations, such as video post-processing and image enhancement.

In some embodiments, media pipeline **316** includes fixed function or programmable logic units to perform one or more specialized media operations, such as video decode acceleration, video de-interlacing, and video encode acceleration in place of, or on behalf of video codec engine **306**. In some embodiments, media pipeline **316** additionally includes a thread spawning unit to spawn threads for execution on 3D/Media sub-system **315**. The spawned threads perform computations for the media operations on one or more graphics execution units included in 3D/Media sub-system **315**. In some embodiments, 3D/Media subsystem **315** includes logic for executing threads spawned by 3D pipeline **312** and media pipeline **316**. In one embodiment, the pipelines send thread execution requests to 3D/Media subsystem **315**, which includes thread dispatch logic for arbitrating and dispatching the various requests to available thread execution resources. The execution resources include an array of graphics execution units to process the 3D and media threads. In some embodiments, 3D/Media subsystem **315** includes one or more internal caches for thread instructions and data. In some embodiments, the subsystem also includes shared memory, including registers and addressable memory, to share data between threads and to store output data.

#### Graphics Processing Engine

FIG. **4** is a block diagram of a graphics processing engine **410** of a graphics processor in accordance with some embodiments. In one embodiment, the graphics processing engine (GPE) **410** is a version of the GPE **310** shown in FIG. **3**. Elements of FIG. **4** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. For example, the 3D pipeline **312** and media pipeline **316** of FIG. **3** are illustrated. The media pipeline **316** is optional in some embodiments of the GPE **410** and may not be explicitly included within the GPE **410**. For example and in at least one embodiment, a separate media and/or image processor is coupled to the GPE **410**.

In some embodiments, GPE **410** couples with or includes a command streamer **403**, which provides a command stream to the 3D pipeline **312** and/or media pipelines **316**. In some embodiments, command streamer **403** is coupled with memory, which can be system memory, or one or more of internal cache memory and shared cache memory. In some embodiments, command streamer **403** receives commands

from the memory and sends the commands to 3D pipeline **312** and/or media pipeline **316**. The commands are directives fetched from a ring buffer, which stores commands for the 3D pipeline **312** and media pipeline **316**. In one embodiment, the ring buffer can additionally include batch command buffers storing batches of multiple commands. The commands for the 3D pipeline **312** can also include references to data stored in memory, such as but not limited to vertex and geometry data for the 3D pipeline **312** and/or image data and memory objects for the media pipeline **316**. The 3D pipeline **312** and media pipeline **316** process the commands and data by performing operations via logic within the respective pipelines or by dispatching one or more execution threads to a graphics core array **414**. In one embodiment the graphics core array **414** include one or more blocks of graphics cores (e.g., graphics core(s) **415A**, graphics core(s) **415B**), each block including one or more graphics cores. Each graphics core includes a set of graphics execution resources that includes general-purpose and graphics specific execution logic to perform graphics and compute operations, as well as fixed function texture processing and/or machine learning and artificial intelligence acceleration logic.

In various embodiments the 3D pipeline **312** includes fixed function and programmable logic to process one or more shader programs, such as vertex shaders, geometry shaders, pixel shaders, fragment shaders, compute shaders, or other shader programs, by processing the instructions and dispatching execution threads to the graphics core array **414**. The graphics core array **414** provides a unified block of execution resources for use in processing these shader programs. Multi-purpose execution logic (e.g., execution units) within the graphics core(s) **415A-414B** of the graphics core array **414** includes support for various 3D API shader languages and can execute multiple simultaneous execution threads associated with multiple shaders.

In some embodiments the graphics core array **414** also includes execution logic to perform media functions, such as video and/or image processing. In one embodiment, the execution units additionally include general-purpose logic that is programmable to perform parallel general-purpose computational operations, in addition to graphics processing operations. The general-purpose logic can perform processing operations in parallel or in conjunction with general-purpose logic within the processor core(s) **107** of FIG. **1** or core **202A-202N** as in FIG. **2**.

Output data generated by threads executing on the graphics core array **414** can output data to memory in a unified return buffer (URB) **418**. The URB **418** can store data for multiple threads. In some embodiments the URB **418** may be used to send data between different threads executing on the graphics core array **414**. In some embodiments the URB **418** may additionally be used for synchronization between threads on the graphics core array and fixed function logic within the shared function logic **420**.

In some embodiments, graphics core array **414** is scalable, such that the array includes a variable number of graphics cores, each having a variable number of execution units based on the target power and performance level of GPE **410**. In one embodiment the execution resources are dynamically scalable, such that execution resources may be enabled or disabled as needed.

The graphics core array **414** couples with shared function logic **420** that includes multiple resources that are shared between the graphics cores in the graphics core array. The shared functions within the shared function logic **420** are hardware logic units that provide specialized supplemental

functionality to the graphics core array **414**. In various embodiments, shared function logic **420** includes but is not limited to sampler **421**, math **422**, and inter-thread communication (ITC) **423** logic. Additionally, some embodiments implement one or more cache(s) **425** within the shared function logic **420**.

A shared function is implemented where the demand for a given specialized function is insufficient for inclusion within the graphics core array **414**. Instead a single instantiation of that specialized function is implemented as a stand-alone entity in the shared function logic **420** and shared among the execution resources within the graphics core array **414**. The precise set of functions that are shared between the graphics core array **414** and included within the graphics core array **414** varies across embodiments. In some embodiments, specific shared functions within the shared function logic **420** that are used extensively by the graphics core array **414** may be included within shared function logic **416** within the graphics core array **414**. In various embodiments, the shared function logic **416** within the graphics core array **414** can include some or all logic within the shared function logic **420**. In one embodiment, all logic elements within the shared function logic **420** may be duplicated within the shared function logic **416** of the graphics core array **414**. In one embodiment the shared function logic **420** is excluded in favor of the shared function logic **416** within the graphics core array **414**.

FIG. **5** is a block diagram of hardware logic of a graphics processor core **500**, according to some embodiments described herein. Elements of FIG. **5** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. The illustrated graphics processor core **500**, in some embodiments, is included within the graphics core array **414** of FIG. **4**. The graphics processor core **500**, sometimes referred to as a core slice, can be one or multiple graphics cores within a modular graphics processor. The graphics processor core **500** is exemplary of one graphics core slice, and a graphics processor as described herein may include multiple graphics core slices based on target power and performance envelopes. Each graphics core **500** can include a fixed function block **530** coupled with multiple sub-cores **501A-501F**, also referred to as sub-slices, that include modular blocks of general-purpose and fixed function logic.

In some embodiments the fixed function block **530** includes a geometry/fixed function pipeline **536** that can be shared by all sub-cores in the graphics processor **500**, for example, in lower performance and/or lower power graphics processor implementations. In various embodiments, the geometry/fixed function pipeline **536** includes a 3D fixed function pipeline (e.g., 3D pipeline **312** as in FIG. **3** and FIG. **4**) a video front-end unit, a thread spawner and thread dispatcher, and a unified return buffer manager, which manages unified return buffers, such as the unified return buffer **418** of FIG. **4**.

In one embodiment the fixed function block **530** also includes a graphics SoC interface **537**, a graphics microcontroller **538**, and a media pipeline **539**. The graphics SoC interface **537** provides an interface between the graphics core **500** and other processor cores within a system on a chip integrated circuit. The graphics microcontroller **538** is a programmable sub-processor that is configurable to manage various functions of the graphics processor **500**, including thread dispatch, scheduling, and pre-emption. The media pipeline **539** (e.g., media pipeline **316** of FIG. **3** and FIG. **4**) includes logic to facilitate the decoding, encoding, pre-

processing, and/or post-processing of multimedia data, including image and video data. The media pipeline **539** implement media operations via requests to compute or sampling logic within the sub-cores **501-501F**.

In one embodiment the SoC interface **537** enables the graphics core **500** to communicate with general-purpose application processor cores (e.g., CPUs) and/or other components within an SoC, including memory hierarchy elements such as a shared last level cache memory, the system RAM, and/or embedded on-chip or on-package DRAM. The SoC interface **537** can also enable communication with fixed function devices within the SoC, such as camera imaging pipelines, and enables the use of and/or implements global memory atomics that may be shared between the graphics core **500** and CPUs within the SoC. The SoC interface **537** can also implement power management controls for the graphics core **500** and enable an interface between a clock domain of the graphic core **500** and other clock domains within the SoC. In one embodiment the SoC interface **537** enables receipt of command buffers from a command streamer and global thread dispatcher that are configured to provide commands and instructions to each of one or more graphics cores within a graphics processor. The commands and instructions can be dispatched to the media pipeline **539**, when media operations are to be performed, or a geometry and fixed function pipeline (e.g., geometry and fixed function pipeline **536**, geometry and fixed function pipeline **514**) when graphics processing operations are to be performed.

The graphics microcontroller **538** can be configured to perform various scheduling and management tasks for the graphics core **500**. In one embodiment the graphics microcontroller **538** can perform graphics and/or compute workload scheduling on the various graphics parallel engines within execution unit (EU) arrays **502A-502F**, **504A-504F** within the sub-cores **501A-501F**. In this scheduling model, host software executing on a CPU core of an SoC including the graphics core **500** can submit workloads one of multiple graphic processor doorbells, which invokes a scheduling operation on the appropriate graphics engine. Scheduling operations include determining which workload to run next, submitting a workload to a command streamer, pre-empting existing workloads running on an engine, monitoring progress of a workload, and notifying host software when a workload is complete. In one embodiment the graphics microcontroller **538** can also facilitate low-power or idle states for the graphics core **500**, providing the graphics core **500** with the ability to save and restore registers within the graphics core **500** across low-power state transitions independently from the operating system and/or graphics driver software on the system.

The graphics core **500** may have greater than or fewer than the illustrated sub-cores **501A-501F**, up to N modular sub-cores. For each set of N sub-cores, the graphics core **500** can also include shared function logic **510**, shared and/or cache memory **512**, a geometry/fixed function pipeline **514**, as well as additional fixed function logic **516** to accelerate various graphics and compute processing operations. The shared function logic **510** can include logic units associated with the shared function logic **420** of FIG. **4** (e.g., sampler, math, and/or inter-thread communication logic) that can be shared by each N sub-cores within the graphics core **500**. The shared and/or cache memory **512** can be a last-level cache for the set of N sub-cores **501A-501F** within the graphics core **500**, and can also serve as shared memory that is accessible by multiple sub-cores. The geometry/fixed function pipeline **514** can be included instead of the geom-

## 11

etry/fixed function pipeline **536** within the fixed function block **530** and can include the same or similar logic units.

In one embodiment the graphics core **500** includes additional fixed function logic **516** that can include various fixed function acceleration logic for use by the graphics core **500**. In one embodiment the additional fixed function logic **516** includes an additional geometry pipeline for use in position only shading. In position-only shading, two geometry pipelines exist, the full geometry pipeline within the geometry/fixed function pipeline **516**, **536**, and a cull pipeline, which is an additional geometry pipeline which may be included within the additional fixed function logic **516**. In one embodiment the cull pipeline is a trimmed down version of the full geometry pipeline. The full pipeline and the cull pipeline can execute different instances of the same application, each instance having a separate context. Position only shading can hide long cull runs of discarded triangles, enabling shading to be completed earlier in some instances. For example and in one embodiment the cull pipeline logic within the additional fixed function logic **516** can execute position shaders in parallel with the main application and generally generates critical results faster than the full pipeline, as the cull pipeline fetches and shades only the position attribute of the vertices, without performing rasterization and rendering of the pixels to the frame buffer. The cull pipeline can use the generated critical results to compute visibility information for all the triangles without regard to whether those triangles are culled. The full pipeline (which in this instance may be referred to as a replay pipeline) can consume the visibility information to skip the culled triangles to shade only the visible triangles that are finally passed to the rasterization phase.

In one embodiment the additional fixed function logic **516** can also include machine-learning acceleration logic, such as fixed function matrix multiplication logic, for implementations including optimizations for machine learning training or inferencing.

Within each graphics sub-core **501A-501F** includes a set of execution resources that may be used to perform graphics, media, and compute operations in response to requests by graphics pipeline, media pipeline, or shader programs. The graphics sub-cores **501A-501F** include multiple EU arrays **502A-502F**, **504A-504F**, thread dispatch and inter-thread communication (TD/IC) logic **503A-503F**, a 3D (e.g., texture) sampler **505A-505F**, a media sampler **506A-506F**, a shader processor **507A-507F**, and shared local memory (SLM) **508A-508F**. The EU arrays **502A-502F**, **504A-504F** each include multiple execution units, which are general-purpose graphics processing units capable of performing floating-point and integer/fixed-point logic operations in service of a graphics, media, or compute operation, including graphics, media, or compute shader programs. The TD/IC logic **503A-503F** performs local thread dispatch and thread control operations for the execution units within a sub-core and facilitate communication between threads executing on the execution units of the sub-core. The 3D sampler **505A-505F** can read texture or other 3D graphics related data into memory. The 3D sampler can read texture data differently based on a configured sample state and the texture format associated with a given texture. The media sampler **506A-506F** can perform similar read operations based on the type and format associated with media data. In one embodiment, each graphics sub-core **501A-501F** can alternately include a unified 3D and media sampler. Threads executing on the execution units within each of the sub-cores **501A-501F** can make use of shared local memory

## 12

**508A-508F** within each sub-core, to enable threads executing within a thread group to execute using a common pool of on-chip memory.

## Execution Units

FIGS. **6A-6B** illustrate thread execution logic **600** including an array of processing elements employed in a graphics processor core according to embodiments described herein. Elements of FIGS. **6A-6B** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. FIG. **6A** illustrates an overview of thread execution logic **600**, which can include a variant of the hardware logic illustrated with each sub-core **501A-501F** of FIG. **5**. FIG. **6B** illustrates exemplary internal details of an execution unit.

As illustrated in FIG. **6A**, in some embodiments thread execution logic **600** includes a shader processor **602**, a thread dispatcher **604**, instruction cache **606**, a scalable execution unit array including a plurality of execution units **608A-608N**, a sampler **610**, a data cache **612**, and a data port **614**. In one embodiment the scalable execution unit array can dynamically scale by enabling or disabling one or more execution units (e.g., any of execution unit **608A**, **608B**, **608C**, **608D**, through **608N-1** and **608N**) based on the computational requirements of a workload. In one embodiment the included components are interconnected via an interconnect fabric that links to each of the components. In some embodiments, thread execution logic **600** includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache **606**, data port **614**, sampler **610**, and execution units **608A-608N**. In some embodiments, each execution unit (e.g. **608A**) is a stand-alone programmable general-purpose computational unit that is capable of executing multiple simultaneous hardware threads while processing multiple data elements in parallel for each thread. In various embodiments, the array of execution units **608A-608N** is scalable to include any number individual execution units.

In some embodiments, the execution units **608A-608N** are primarily used to execute shader programs. A shader processor **602** can process the various shader programs and dispatch execution threads associated with the shader programs via a thread dispatcher **604**. In one embodiment the thread dispatcher includes logic to arbitrate thread initiation requests from the graphics and media pipelines and instantiate the requested threads on one or more execution unit in the execution units **608A-608N**. For example, a geometry pipeline can dispatch vertex, tessellation, or geometry shaders to the thread execution logic for processing. In some embodiments, thread dispatcher **604** can also process runtime thread spawning requests from the executing shader programs.

In some embodiments, the execution units **608A-608N** support an instruction set that includes native support for many standard 3D graphics shader instructions, such that shader programs from graphics libraries (e.g., Direct 3D and OpenGL) are executed with a minimal translation. The execution units support vertex and geometry processing (e.g., vertex programs, geometry programs, vertex shaders), pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders). Each of the execution units **608A-608N** is capable of multi-issue single instruction multiple data (SIMD) execution and multi-threaded operation enables an efficient execution environment in the face of higher latency memory accesses. Each hardware thread within each execution unit has a dedicated high-bandwidth register file and associated

independent thread-state. Execution is multi-issue per clock to pipelines capable of integer, single and double precision floating point operations, SIMD branch capability, logical operations, transcendental operations, and other miscellaneous operations. While waiting for data from memory or one of the shared functions, dependency logic within the execution units **608A-608N** causes a waiting thread to sleep until the requested data has been returned. While the waiting thread is sleeping, hardware resources may be devoted to processing other threads. For example, during a delay associated with a vertex shader operation, an execution unit can perform operations for a pixel shader, fragment shader, or another type of shader program, including a different vertex shader.

Each execution unit in execution units **608A-608N** operates on arrays of data elements. The number of data elements is the “execution size,” or the number of channels for the instruction. An execution channel is a logical unit of execution for data element access, masking, and flow control within instructions. The number of channels may be independent of the number of physical Arithmetic Logic Units (ALUs) or Floating Point Units (FPUs) for a particular graphics processor. In some embodiments, execution units **608A-608N** support integer and floating-point data types.

The execution unit instruction set includes SIMD instructions. The various data elements can be stored as a packed data type in a register and the execution unit will process the various elements based on the data size of the elements. For example, when operating on a 256-bit wide vector, the 256 bits of the vector are stored in a register and the execution unit operates on the vector as four separate 64-bit packed data elements (Quad-Word (QW) size data elements), eight separate 32-bit packed data elements (Double Word (DW) size data elements), sixteen separate 16-bit packed data elements (Word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). However, different vector widths and register sizes are possible.

In one embodiment one or more execution units can be combined into a fused execution unit **609A-609N** having thread control logic (**607A-607N**) that is common to the fused EUs. Multiple EUs can be fused into an EU group. Each EU in the fused EU group can be configured to execute a separate SIMD hardware thread. The number of EUs in a fused EU group can vary according to embodiments. Additionally, various SIMD widths can be performed per-EU, including but not limited to SIMD8, SIMD16, and SIMD32. Each fused graphics execution unit **609A-609N** includes at least two execution units. For example, fused execution unit **609A** includes a first EU **608A**, second EU **608B**, and thread control logic **607A** that is common to the first EU **608A** and the second EU **608B**. The thread control logic **607A** controls threads executed on the fused graphics execution unit **609A**, allowing each EU within the fused execution units **609A-609N** to execute using a common instruction pointer register.

One or more internal instruction caches (e.g., **606**) are included in the thread execution logic **600** to cache thread instructions for the execution units. In some embodiments, one or more data caches (e.g., **612**) are included to cache thread data during thread execution. In some embodiments, a sampler **610** is included to provide texture sampling for 3D operations and media sampling for media operations. In some embodiments, sampler **610** includes specialized texture or media sampling functionality to process texture or media data during the sampling process before providing the sampled data to an execution unit.

During execution, the graphics and media pipelines send thread initiation requests to thread execution logic **600** via thread spawning and dispatch logic. Once a group of geometric objects has been processed and rasterized into pixel data, pixel processor logic (e.g., pixel shader logic, fragment shader logic, etc.) within the shader processor **602** is invoked to further compute output information and cause results to be written to output surfaces (e.g., color buffers, depth buffers, stencil buffers, etc.). In some embodiments, a pixel shader or fragment shader calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some embodiments, pixel processor logic within the shader processor **602** then executes an application programming interface (API)-supplied pixel or fragment shader program. To execute the shader program, the shader processor **602** dispatches threads to an execution unit (e.g., **608A**) via thread dispatcher **604**. In some embodiments, shader processor **602** uses texture sampling logic in the sampler **610** to access texture data in texture maps stored in memory. Arithmetic operations on the texture data and the input geometry data compute pixel color data for each geometric fragment, or discards one or more pixels from further processing.

In some embodiments, the data port **614** provides a memory access mechanism for the thread execution logic **600** to output processed data to memory for further processing on a graphics processor output pipeline. In some embodiments, the data port **614** includes or couples to one or more cache memories (e.g., data cache **612**) to cache data for memory access via the data port.

As illustrated in FIG. 6B, a graphics execution unit **608** can include an instruction fetch unit **637**, a general register file array (GRF) **624**, an architectural register file array (ARF) **626**, a thread arbiter **622**, a send unit **630**, a branch unit **632**, a set of SIMD floating point units (FPUs) **634**, and in one embodiment a set of dedicated integer SIMD ALUs **635**. The GRF **624** and ARF **626** includes the set of general register files and architecture register files associated with each simultaneous hardware thread that may be active in the graphics execution unit **608**. In one embodiment, per thread architectural state is maintained in the ARF **626**, while data used during thread execution is stored in the GRF **624**. The execution state of each thread, including the instruction pointers for each thread, can be held in thread-specific registers in the ARF **626**.

In one embodiment the graphics execution unit **608** has an architecture that is a combination of Simultaneous Multi-Threading (SMT) and fine-grained Interleaved Multi-Threading (IMT). The architecture has a modular configuration that can be fine-tuned at design time based on a target number of simultaneous threads and number of registers per execution unit, where execution unit resources are divided across logic used to execute multiple simultaneous threads.

In one embodiment, the graphics execution unit **608** can co-issue multiple instructions, which may each be different instructions. The thread arbiter **622** of the graphics execution unit thread **608** can dispatch the instructions to one of the send unit **630**, branch unit **642**, or SIMD FPU(s) **634** for execution. Each execution thread can access 128 general-purpose registers within the GRF **624**, where each register can store 32 bytes, accessible as a SIMD 8-element vector of 32-bit data elements. In one embodiment, each execution unit thread has access to 4 Kbytes within the GRF **624**, although embodiments are not so limited, and greater or fewer register resources may be provided in other embodiments. In one embodiment up to seven threads can execute simultaneously, although the number of threads per execu-

## 15

tion unit can also vary according to embodiments. In an embodiment in which seven threads may access 4 Kbytes, the GRF **624** can store a total of 28 Kbytes. Flexible addressing modes can permit registers to be addressed together to build effectively wider registers or to represent strided rectangular block data structures.

In one embodiment, memory operations, sampler operations, and other longer-latency system communications are dispatched via “send” instructions that are executed by the message passing send unit **630**. In one embodiment, branch instructions are dispatched to a dedicated branch unit **632** to facilitate SIMD divergence and eventual convergence.

In one embodiment the graphics execution unit **608** includes one or more SIMD floating point units (FPU(s)) **634** to perform floating-point operations. In one embodiment, the FPU(s) **634** also support integer computation. In one embodiment the FPU(s) **634** can SIMD execute up to M number of 32-bit floating-point (or integer) operations, or SIMD execute up to 2M 16-bit integer or 16-bit floating-point operations.

In one embodiment, at least one of the FPU(s) provides extended math capability to support high-throughput transcendental math functions and double precision 64-bit floating-point. In some embodiments, a set of 8-bit integer SIMD ALUs **635** are also present, and may be specifically optimized to perform operations associated with machine learning computations.

In one embodiment, arrays of multiple instances of the graphics execution unit **608** can be instantiated in a graphics sub-core grouping (e.g., a sub-slice). For scalability, product architects can choose the exact number of execution units per sub-core grouping. In one embodiment the execution unit **608** can execute instructions across a plurality of execution channels. In a further embodiment, each thread executed on the graphics execution unit **608** is executed on a different channel.

FIG. **7** is a block diagram illustrating a graphics processor instruction formats **700** according to some embodiments. In one or more embodiment, the graphics processor execution units support an instruction set having instructions in multiple formats. The solid lined boxes illustrate the components that are generally included in an execution unit instruction, while the dashed lines include components that are optional or that are only included in a sub-set of the instructions. In some embodiments, instruction format **700** described and illustrated are macro-instructions, in that they are instructions supplied to the execution unit, as opposed to micro-operations resulting from instruction decode once the instruction is processed.

In some embodiments, the graphics processor execution units natively support instructions in a 128-bit instruction format **710**. A 64-bit compacted instruction format **730** is available for some instructions based on the selected instruction, instruction options, and number of operands. The native 128-bit instruction format **710** provides access to all instruction options, while some options and operations are restricted in the 64-bit format **730**. The native instructions available in the 64-bit format **730** vary by embodiment. In some embodiments, the instruction is compacted in part using a set of index values in an index field **713**. The execution unit hardware references a set of compaction tables based on the index values and uses the compaction table outputs to reconstruct a native instruction in the 128-bit instruction format **710**.

For each format, instruction opcode **712** defines the operation that the execution unit is to perform. The execution units execute each instruction in parallel across the

## 16

multiple data elements of each operand. For example, in response to an add instruction the execution unit performs a simultaneous add operation across each color channel representing a texture element or picture element. By default, the execution unit performs each instruction across all data channels of the operands. In some embodiments, instruction control field **714** enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For instructions in the 128-bit instruction format **710** an exec-size field **716** limits the number of data channels that will be executed in parallel. In some embodiments, exec-size field **716** is not available for use in the 64-bit compact instruction format **730**. Some execution unit instructions have up to three operands including two source operands, src0 **720**, src1 **722**, and one destination **718**. In some embodiments, the execution units support dual destination instructions, where one of the destinations is implied. Data manipulation instructions can have a third source operand (e.g., SRC2 **724**), where the instruction opcode **712** determines the number of source operands. An instruction's last source operand can be an immediate (e.g., hard-coded) value passed with the instruction.

In some embodiments, the 128-bit instruction format **710** includes an access/address mode field **726** specifying, for example, whether direct register addressing mode or indirect register addressing mode is used. When direct register addressing mode is used, the register address of one or more operands is directly provided by bits in the instruction.

In some embodiments, the 128-bit instruction format **710** includes an access/address mode field **726**, which specifies an address mode and/or an access mode for the instruction. In one embodiment the access mode is used to define a data access alignment for the instruction. Some embodiments support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in a first mode, the instruction may use byte-aligned addressing for source and destination operands and when in a second mode, the instruction may use 16-byte-aligned addressing for all source and destination operands.

In one embodiment, the address mode portion of the access/address mode field **726** determines whether the instruction is to use direct or indirect addressing. When direct register addressing mode is used bits in the instruction directly provide the register address of one or more operands. When indirect register addressing mode is used, the register address of one or more operands may be computed based on an address register value and an address immediate field in the instruction.

In some embodiments instructions are grouped based on opcode **712** bit-fields to simplify Opcode decode **740**. For an 8-bit opcode, bits **4**, **5**, and **6** allow the execution unit to determine the type of opcode. The precise opcode grouping shown is merely an example. In some embodiments, a move and logic opcode group **742** includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some embodiments, move and logic group **742** shares the five most significant bits (MSB), where move (mov) instructions are in the form of 0000xxxxb and logic instructions are in the form of 0001xxxxb. A flow control instruction group **744** (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group **746** includes a mix of instructions, including synchronization instructions (e.g., wait, send) in the form of 0011xxxxb (e.g., 0x30). A parallel math instruction group

**748** includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group **748** performs the arithmetic operations in parallel across data channels. The vector math group **750** includes arithmetic instructions (e.g., dp4) in the form of 0101xxxxb (e.g., 0x50). The vector math group performs arithmetic such as dot product calculations on vector operands.

#### Graphics Pipeline

FIG. **8** is a block diagram of another embodiment of a graphics processor **800**. Elements of FIG. **8** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

In some embodiments, graphics processor **800** includes a geometry pipeline **820**, a media pipeline **830**, a display engine **840**, thread execution logic **850**, and a render output pipeline **870**. In some embodiments, graphics processor **800** is a graphics processor within a multi-core processing system that includes one or more general-purpose processing cores. The graphics processor is controlled by register writes to one or more control registers (not shown) or via commands issued to graphics processor **800** via a ring interconnect **802**. In some embodiments, ring interconnect **802** couples graphics processor **800** to other processing components, such as other graphics processors or general-purpose processors. Commands from ring interconnect **802** are interpreted by a command streamer **803**, which supplies instructions to individual components of the geometry pipeline **820** or the media pipeline **830**.

In some embodiments, command streamer **803** directs the operation of a vertex fetcher **805** that reads vertex data from memory and executes vertex-processing commands provided by command streamer **803**. In some embodiments, vertex fetcher **805** provides vertex data to a vertex shader **807**, which performs coordinate space transformation and lighting operations to each vertex. In some embodiments, vertex fetcher **805** and vertex shader **807** execute vertex-processing instructions by dispatching execution threads to execution units **852A-852B** via a thread dispatcher **831**.

In some embodiments, execution units **852A-852B** are an array of vector processors having an instruction set for performing graphics and media operations. In some embodiments, execution units **852A-852B** have an attached L1 cache **851** that is specific for each array or shared between the arrays. The cache can be configured as a data cache, an instruction cache, or a single cache that is partitioned to contain data and instructions in different partitions.

In some embodiments, geometry pipeline **820** includes tessellation components to perform hardware-accelerated tessellation of 3D objects. In some embodiments, a programmable hull shader **811** configures the tessellation operations. A programmable domain shader **817** provides back-end evaluation of tessellation output. A tessellator **813** operates at the direction of hull shader **811** and contains special purpose logic to generate a set of detailed geometric objects based on a coarse geometric model that is provided as input to geometry pipeline **820**. In some embodiments, if tessellation is not used, tessellation components (e.g., hull shader **811**, tessellator **813**, and domain shader **817**) can be bypassed.

In some embodiments, complete geometric objects can be processed by a geometry shader **819** via one or more threads dispatched to execution units **852A-852B**, or can proceed directly to the clipper **829**. In some embodiments, the geometry shader operates on entire geometric objects, rather

than vertices or patches of vertices as in previous stages of the graphics pipeline. If the tessellation is disabled the geometry shader **819** receives input from the vertex shader **807**. In some embodiments, geometry shader **819** is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

Before rasterization, a clipper **829** processes vertex data. The clipper **829** may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some embodiments, a rasterizer and depth test component **873** in the render output pipeline **870** dispatches pixel shaders to convert the geometric objects into per pixel representations. In some embodiments, pixel shader logic is included in thread execution logic **850**. In some embodiments, an application can bypass the rasterizer and depth test component **873** and access un-rasterized vertex data via a stream out unit **823**. The graphics processor **800** has an interconnect bus, interconnect fabric, or some other interconnect mechanism that allows data and message passing amongst the major components of the processor. In some embodiments, execution units **852A-852B** and associated logic units (e.g., L1 cache **851**, sampler **854**, texture cache **858**, etc.) interconnect via a data port **856** to perform memory access and communicate with render output pipeline components of the processor. In some embodiments, sampler **854**, caches **851**, **858** and execution units **852A-852B** each have separate memory access paths. In one embodiment the texture cache **858** can also be configured as a sampler cache. In some embodiments, render output pipeline **870** contains a rasterizer and depth test component **873** that converts vertex-based objects into an associated pixel-based representation. In some embodiments, the rasterizer logic includes a windower/masker unit to perform fixed function triangle and line rasterization. An associated render cache **878** and depth cache **879** are also available in some embodiments. A pixel operations component **877** performs pixel-based operations on the data, though in some instances, pixel operations associated with 2D operations (e.g. bit block image transfers with blending) are performed by the 2D engine **841**, or substituted at display time by the display controller **843** using overlay display planes. In some embodiments, a shared L3 cache **875** is available to all graphics components, allowing the sharing of data without the use of main system memory.

In some embodiments, graphics processor media pipeline **830** includes a media engine **837** and a video front-end **834**. In some embodiments, video front-end **834** receives pipeline commands from the command streamer **803**. In some embodiments, media pipeline **830** includes a separate command streamer. In some embodiments, video front-end **834** processes media commands before sending the command to the media engine **837**. In some embodiments, media engine **837** includes thread spawning functionality to spawn threads for dispatch to thread execution logic **850** via thread dispatcher **831**.

In some embodiments, graphics processor **800** includes a display engine **840**. In some embodiments, display engine **840** is external to processor **800** and couples with the graphics processor via the ring interconnect **802**, or some other interconnect bus or fabric. In some embodiments, display engine **840** includes a 2D engine **841** and a display controller **843**. In some embodiments, display engine **840** contains special purpose logic capable of operating independently of the 3D pipeline. In some embodiments, display controller **843** couples with a display device (not shown),

which may be a system integrated display device, as in a laptop computer, or an external display device attached via a display device connector.

In some embodiments, the geometry pipeline **820** and media pipeline **830** are configurable to perform operations based on multiple graphics and media programming interfaces and are not specific to any one application programming interface (API). In some embodiments, driver software for the graphics processor translates API calls that are specific to a particular graphics or media library into commands that can be processed by the graphics processor. In some embodiments, support is provided for the Open Graphics Library (OpenGL), Open Computing Language (OpenCL), and/or Vulkan graphics and compute API, all from the Khronos Group. In some embodiments, support may also be provided for the Direct3D library from the Microsoft Corporation. In some embodiments, a combination of these libraries may be supported. Support may also be provided for the Open Source Computer Vision Library (OpenCV). A future API with a compatible 3D pipeline would also be supported if a mapping can be made from the pipeline of the future API to the pipeline of the graphics processor.

#### Graphics Pipeline Programming

FIG. 9A is a block diagram illustrating a graphics processor command format **900** according to some embodiments. FIG. 9B is a block diagram illustrating a graphics processor command sequence **910** according to an embodiment. The solid lined boxes in FIG. 9A illustrate the components that are generally included in a graphics command while the dashed lines include components that are optional or that are only included in a sub-set of the graphics commands. The exemplary graphics processor command format **900** of FIG. 9A includes data fields to identify a client **902**, a command operation code (opcode) **904**, and data **906** for the command. A sub-opcode **905** and a command size **908** are also included in some commands.

In some embodiments, client **902** specifies the client unit of the graphics device that processes the command data. In some embodiments, a graphics processor command parser examines the client field of each command to condition the further processing of the command and route the command data to the appropriate client unit. In some embodiments, the graphics processor client units include a memory interface unit, a render unit, a 2D unit, a 3D unit, and a media unit. Each client unit has a corresponding processing pipeline that processes the commands. Once the command is received by the client unit, the client unit reads the opcode **904** and, if present, sub-opcode **905** to determine the operation to perform. The client unit performs the command using information in data field **906**. For some commands an explicit command size **908** is expected to specify the size of the command. In some embodiments, the command parser automatically determines the size of at least some of the commands based on the command opcode. In some embodiments commands are aligned via multiples of a double word.

The flow diagram in FIG. 9B illustrates an exemplary graphics processor command sequence **910**. In some embodiments, software or firmware of a data processing system that features an embodiment of a graphics processor uses a version of the command sequence shown to set up, execute, and terminate a set of graphics operations. A sample command sequence is shown and described for purposes of example only as embodiments are not limited to these specific commands or to this command sequence. Moreover, the commands may be issued as batch of commands in a

command sequence, such that the graphics processor will process the sequence of commands in at least partially concurrence.

In some embodiments, the graphics processor command sequence **910** may begin with a pipeline flush command **912** to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some embodiments, the 3D pipeline **922** and the media pipeline **924** do not operate concurrently. The pipeline flush is performed to cause the active graphics pipeline to complete any pending commands. In response to a pipeline flush, the command parser for the graphics processor will pause command processing until the active drawing engines complete pending operations and the relevant read caches are invalidated. Optionally, any data in the render cache that is marked 'dirty' can be flushed to memory. In some embodiments, pipeline flush command **912** can be used for pipeline synchronization or before placing the graphics processor into a low power state.

In some embodiments, a pipeline select command **913** is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command **913** is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some embodiments, a pipeline flush command **912** is required immediately before a pipeline switch via the pipeline select command **913**.

In some embodiments, a pipeline control command **914** configures a graphics pipeline for operation and is used to program the 3D pipeline **922** and the media pipeline **924**. In some embodiments, pipeline control command **914** configures the pipeline state for the active pipeline. In one embodiment, the pipeline control command **914** is used for pipeline synchronization and to clear data from one or more cache memories within the active pipeline before processing a batch of commands.

In some embodiments, return buffer state commands **916** are used to configure a set of return buffers for the respective pipelines to write data. Some pipeline operations require the allocation, selection, or configuration of one or more return buffers into which the operations write intermediate data during processing. In some embodiments, the graphics processor also uses one or more return buffers to store output data and to perform cross thread communication. In some embodiments, the return buffer state **916** includes selecting the size and number of return buffers to use for a set of pipeline operations.

The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination **920**, the command sequence is tailored to the 3D pipeline **922** beginning with the 3D pipeline state **930** or the media pipeline **924** beginning at the media pipeline state **940**.

The commands to configure the 3D pipeline state **930** include 3D state setting commands for vertex buffer state, vertex element state, constant color state, depth buffer state, and other state variables that are to be configured before 3D primitive commands are processed. The values of these commands are determined at least in part based on the particular 3D API in use. In some embodiments, 3D pipeline state **930** commands are also able to selectively disable or bypass certain pipeline elements if those elements will not be used.

In some embodiments, 3D primitive **932** command is used to submit 3D primitives to be processed by the 3D pipeline. Commands and associated parameters that are

passed to the graphics processor via the 3D primitive **932** command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D primitive **932** command data to generate vertex data structures. The vertex data structures are stored in one or more return buffers. In some embodiments, 3D primitive **932** command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline **922** dispatches shader execution threads to graphics processor execution units.

In some embodiments, 3D pipeline **922** is triggered via an execute **934** command or event. In some embodiments, a register write triggers command execution. In some embodiments execution is triggered via a 'go' or 'kick' command in the command sequence. In one embodiment, command execution is triggered using a pipeline synchronization command to flush the command sequence through the graphics pipeline. The 3D pipeline will perform geometry processing for the 3D primitives. Once operations are complete, the resulting geometric objects are rasterized and the pixel engine colors the resulting pixels. Additional commands to control pixel shading and pixel back end operations may also be included for those operations.

In some embodiments, the graphics processor command sequence **910** follows the media pipeline **924** path when performing media operations. In general, the specific use and manner of programming for the media pipeline **924** depends on the media or compute operations to be performed. Specific media decode operations may be offloaded to the media pipeline during media decode. In some embodiments, the media pipeline can also be bypassed and media decode can be performed in whole or in part using resources provided by one or more general-purpose processing cores. In one embodiment, the media pipeline also includes elements for general-purpose graphics processor unit (GPGPU) operations, where the graphics processor is used to perform SIMD vector operations using computational shader programs that are not explicitly related to the rendering of graphics primitives.

In some embodiments, media pipeline **924** is configured in a similar manner as the 3D pipeline **922**. A set of commands to configure the media pipeline state **940** are dispatched or placed into a command queue before the media object commands **942**. In some embodiments, commands for the media pipeline state **940** include data to configure the media pipeline elements that will be used to process the media objects. This includes data to configure the video decode and video encode logic within the media pipeline, such as encode or decode format. In some embodiments, commands for the media pipeline state **940** also support the use of one or more pointers to "indirect" state elements that contain a batch of state settings.

In some embodiments, media object commands **942** supply pointers to media objects for processing by the media pipeline. The media objects include memory buffers containing video data to be processed. In some embodiments, all media pipeline states must be valid before issuing a media object command **942**. Once the pipeline state is configured and media object commands **942** are queued, the media pipeline **924** is triggered via an execute command **944** or an equivalent execute event (e.g., register write). Output from media pipeline **924** may then be post processed by operations provided by the 3D pipeline **922** or the media pipeline **924**. In some embodiments, GPGPU operations are configured and executed in a similar manner as media operations.

## Graphics Software Architecture

FIG. 10 illustrates exemplary graphics software architecture for a data processing system **1000** according to some embodiments. In some embodiments, software architecture includes a 3D graphics application **1010**, an operating system **1020**, and at least one processor **1030**. In some embodiments, processor **1030** includes a graphics processor **1032** and one or more general-purpose processor core(s) **1034**. The graphics application **1010** and operating system **1020** each execute in the system memory **1050** of the data processing system.

In some embodiments, 3D graphics application **1010** contains one or more shader programs including shader instructions **1012**. The shader language instructions may be in a high-level shader language, such as the High Level Shader Language (HLSL) or the OpenGL Shader Language (GLSL). The application also includes executable instructions **1014** in a machine language suitable for execution by the general-purpose processor core **1034**. The application also includes graphics objects **1016** defined by vertex data.

In some embodiments, operating system **1020** is a Microsoft® Windows® operating system from the Microsoft Corporation, a proprietary UNIX-like operating system, or an open source UNIX-like operating system using a variant of the Linux kernel. The operating system **1020** can support a graphics API **1022** such as the Direct3D API, the OpenGL API, or the Vulkan API. When the Direct3D API is in use, the operating system **1020** uses a front-end shader compiler **1024** to compile any shader instructions **1012** in HLSL into a lower-level shader language. The compilation may be a just-in-time (JIT) compilation or the application can perform shader pre-compilation. In some embodiments, high-level shaders are compiled into low-level shaders during the compilation of the 3D graphics application **1010**. In some embodiments, the shader instructions **1012** are provided in an intermediate form, such as a version of the Standard Portable Intermediate Representation (SPIR) used by the Vulkan API.

In some embodiments, user mode graphics driver **1026** contains a back-end shader compiler **1027** to convert the shader instructions **1012** into a hardware specific representation. When the OpenGL API is in use, shader instructions **1012** in the GLSL high-level language are passed to a user mode graphics driver **1026** for compilation. In some embodiments, user mode graphics driver **1026** uses operating system kernel mode functions **1028** to communicate with a kernel mode graphics driver **1029**. In some embodiments, kernel mode graphics driver **1029** communicates with graphics processor **1032** to dispatch commands and instructions.

## IP Core Implementations

One or more aspects of at least one embodiment may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as "IP cores," are reusable units of logic for an integrated circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated

such that the circuit performs operations described in association with any of the embodiments described herein.

FIG. 11A is a block diagram illustrating an IP core development system 1100 that may be used to manufacture an integrated circuit to perform operations according to an embodiment. The IP core development system 1100 may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SOC integrated circuit). A design facility 1130 can generate a software simulation 1110 of an IP core design in a high-level programming language (e.g., C/C++). The software simulation 1110 can be used to design, test, and verify the behavior of the IP core using a simulation model 1112. The simulation model 1112 may include functional, behavioral, and/or timing simulations. A register transfer level (RTL) design 1115 can then be created or synthesized from the simulation model 1112. The RTL design 1115 is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design 1115, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

The RTL design 1115 or equivalent may be further synthesized by the design facility into a hardware model 1120, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3<sup>rd</sup> party fabrication facility 1165 using non-volatile memory 1140 (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection 1150 or wireless connection 1160. The fabrication facility 1165 may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

FIG. 11B illustrates a cross-section side view of an integrated circuit package assembly 1170, according to some embodiments described herein. The integrated circuit package assembly 1170 illustrates an implementation of one or more processor or accelerator devices as described herein. The package assembly 1170 includes multiple units of hardware logic 1172, 1174 connected to a substrate 1180. The logic 1172, 1174 may be implemented at least partly in configurable logic or fixed-functionality logic hardware, and can include one or more portions of any of the processor core(s), graphics processor(s), or other accelerator devices described herein. Each unit of logic 1172, 1174 can be implemented within a semiconductor die and coupled with the substrate 1180 via an interconnect structure 1173. The interconnect structure 1173 may be configured to route electrical signals between the logic 1172, 1174 and the substrate 1180, and can include interconnects such as, but not limited to bumps or pillars. In some embodiments, the interconnect structure 1173 may be configured to route electrical signals such as, for example, input/output (I/O) signals and/or power or ground signals associated with the operation of the logic 1172, 1174. In some embodiments, the substrate 1180 is an epoxy-based laminate substrate. The package substrate 1180 may include other suitable types of substrates in other embodiments. The package assembly 1170 can be connected to other electrical devices via a

package interconnect 1183. The package interconnect 1183 may be coupled to a surface of the substrate 1180 to route electrical signals to other electrical devices, such as a motherboard, other chipset, or multi-chip module.

In some embodiments, the units of logic 1172, 1174 are electrically coupled with a bridge 1182 that is configured to route electrical signals between the logic 1172, 1174. The bridge 1182 may be a dense interconnect structure that provides a route for electrical signals. The bridge 1182 may include a bridge substrate composed of glass or a suitable semiconductor material. Electrical routing features can be formed on the bridge substrate to provide a chip-to-chip connection between the logic 1172, 1174.

Although two units of logic 1172, 1174 and a bridge 1182 are illustrated, embodiments described herein may include more or fewer logic units on one or more dies. The one or more dies may be connected by zero or more bridges, as the bridge 1182 may be excluded when the logic is included on a single die. Alternatively, multiple dies or units of logic can be connected by one or more bridges. Additionally, multiple logic units, dies, and bridges can be connected together in other possible configurations, including three-dimensional configurations.

#### 25 Exemplary System on a Chip Integrated Circuit

FIGS. 12-14 illustrated exemplary integrated circuits and associated graphics processors that may be fabricated using one or more IP cores, according to various embodiments described herein. In addition to what is illustrated, other logic and circuits may be included, including additional graphics processors/cores, peripheral interface controllers, or general-purpose processor cores.

FIG. 12 is a block diagram illustrating an exemplary system on a chip integrated circuit 1200 that may be fabricated using one or more IP cores, according to an embodiment. Exemplary integrated circuit 1200 includes one or more application processor(s) 1205 (e.g., CPUs), at least one graphics processor 1210, and may additionally include an image processor 1215 and/or a video processor 1220, any of which may be a modular IP core from the same or multiple different design facilities. Integrated circuit 1200 includes peripheral or bus logic including a USB controller 1225, UART controller 1230, an SPI/SDIO controller 1235, and an I<sup>2</sup>S/I<sup>2</sup>C controller 1240. Additionally, the integrated circuit can include a display device 1245 coupled to one or more of a high-definition multimedia interface (HDMI) controller 1250 and a mobile industry processor interface (MIPI) display interface 1255. Storage may be provided by a flash memory subsystem 1260 including flash memory and a flash memory controller. Memory interface may be provided via a memory controller 1265 for access to SDRAM or SRAM memory devices. Some integrated circuits additionally include an embedded security engine 1270.

FIGS. 13A-13B are block diagrams illustrating exemplary graphics processors for use within an SoC, according to embodiments described herein. FIG. 13A illustrates an exemplary graphics processor 1310 of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. FIG. 13B illustrates an additional exemplary graphics processor 1340 of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. Graphics processor 1310 of FIG. 13A is an example of a low power graphics processor core. Graphics processor 1340 of FIG. 13B is an example of a higher performance graphics processor core. Each of the graphics processors 1310, 1340 can be variants of the graphics processor 1210 of FIG. 12.

## 25

As shown in FIG. 13A, graphics processor 1310 includes a vertex processor 1305 and one or more fragment processor(s) 1315A-1315N (e.g., 1315A, 1315B, 1315C, 1315D, through 1315N-1, and 1315N). Graphics processor 1310 can execute different shader programs via separate logic, such that the vertex processor 1305 is optimized to execute operations for vertex shader programs, while the one or more fragment processor(s) 1315A-1315N execute fragment (e.g., pixel) shading operations for fragment or pixel shader programs. The vertex processor 1305 performs the vertex processing stage of the 3D graphics pipeline and generates primitives and vertex data. The fragment processor(s) 1315A-1315N use the primitive and vertex data generated by the vertex processor 1305 to produce a frame-buffer that is displayed on a display device. In one embodiment, the fragment processor(s) 1315A-1315N are optimized to execute fragment shader programs as provided for in the OpenGL API, which may be used to perform similar operations as a pixel shader program as provided for in the Direct 3D API.

Graphics processor 1310 additionally includes one or more memory management units (MMUs) 1320A-1320B, cache(s) 1325A-1325B, and circuit interconnect(s) 1330A-1330B. The one or more MMU(s) 1320A-1320B provide for virtual to physical address mapping for the graphics processor 1310, including for the vertex processor 1305 and/or fragment processor(s) 1315A-1315N, which may reference vertex or image/texture data stored in memory, in addition to vertex or image/texture data stored in the one or more cache(s) 1325A-1325B. In one embodiment the one or more MMU(s) 1320A-1320B may be synchronized with other MMUs within the system, including one or more MMUs associated with the one or more application processor(s) 1205, image processor 1215, and/or video processor 1220 of FIG. 12, such that each processor 1205-1220 can participate in a shared or unified virtual memory system. The one or more circuit interconnect(s) 1330A-1330B enable graphics processor 1310 to interface with other IP cores within the SoC, either via an internal bus of the SoC or via a direct connection, according to embodiments.

As shown FIG. 13B, graphics processor 1340 includes the one or more MMU(s) 1320A-1320B, caches 1325A-1325B, and circuit interconnects 1330A-1330B of the graphics processor 1310 of FIG. 13A. Graphics processor 1340 includes one or more shader core(s) 1355A-1355N (e.g., 1355A, 1355B, 1355C, 1355D, 1355E, 1355F, through 1355N-1, and 1355N), which provides for a unified shader core architecture in which a single core or type or core can execute all types of programmable shader code, including shader program code to implement vertex shaders, fragment shaders, and/or compute shaders. The exact number of shader cores present can vary among embodiments and implementations. Additionally, graphics processor 1340 includes an inter-core task manager 1345, which acts as a thread dispatcher to dispatch execution threads to one or more shader cores 1355A-1355N and a tiling unit 1358 to accelerate tiling operations for tile-based rendering, in which rendering operations for a scene are subdivided in image space, for example to exploit local spatial coherence within a scene or to optimize use of internal caches.

FIGS. 14A-14B illustrate additional exemplary graphics processor logic according to embodiments described herein. FIG. 14A illustrates a graphics core 1400 that may be included within the graphics processor 1210 of FIG. 12, and may be a unified shader core 1355A-1355N as in FIG. 13B.

## 26

FIG. 14B illustrates a highly-parallel general-purpose graphics processing unit 1430 suitable for deployment on a multi-chip module.

As shown in FIG. 14A, the graphics core 1400 includes a shared instruction cache 1402, a texture unit 1418, and a cache/shared memory 1420 that are common to the execution resources within the graphics core 1400. The graphics core 1400 can include multiple slices 1401A-1401N or partition for each core, and a graphics processor can include multiple instances of the graphics core 1400. The slices 1401A-1401N can include support logic including a local instruction cache 1404A-1404N, a thread scheduler 1406A-1406N, a thread dispatcher 1408A-1408N, and a set of registers 1410A. To perform logic operations, the slices 1401A-1401N can include a set of additional function units (AFUs 1412A-1412N), floating-point units (FPU 1414A-1414N), integer arithmetic logic units (ALUs 1416-1416N), address computational units (ACU 1413A-1413N), double-precision floating-point units (DPFPU 1415A-1415N), and matrix processing units (MPU 1417A-1417N).

Some of the computational units operate at a specific precision. For example, the FPUs 1414A-1414N can perform single-precision (32-bit) and half-precision (16-bit) floating point operations, while the DPFPU 1415A-1415N perform double precision (64-bit) floating point operations. The ALUs 1416A-1416N can perform variable precision integer operations at 8-bit, 16-bit, and 32-bit precision, and can be configured for mixed precision operations. The MPUs 1417A-1417N can also be configured for mixed precision matrix operations, including half-precision floating point and 8-bit integer operations. The MPUs 1417-1417N can perform a variety of matrix operations to accelerate machine learning application frameworks, including enabling support for accelerated general matrix to matrix multiplication (GEMM). The AFUs 1412A-1412N can perform additional logic operations not supported by the floating-point or integer units, including trigonometric operations (e.g., Sine, Cosine, etc.).

As shown in FIG. 14B, a general-purpose processing unit (GPGPU) 1430 can be configured to enable highly-parallel compute operations to be performed by an array of graphics processing units. Additionally, the GPGPU 1430 can be linked directly to other instances of the GPGPU to create a multi-GPU cluster to improve training speed for particularly deep neural networks. The GPGPU 1430 includes a host interface 1432 to enable a connection with a host processor. In one embodiment the host interface 1432 is a PCI Express interface. However, the host interface can also be a vendor specific communications interface or communications fabric. The GPGPU 1430 receives commands from the host processor and uses a global scheduler 1434 to distribute execution threads associated with those commands to a set of compute clusters 1436A-1436H. The compute clusters 1436A-1436H share a cache memory 1438. The cache memory 1438 can serve as a higher-level cache for cache memories within the compute clusters 1436A-1436H.

The GPGPU 1430 includes memory 1434A-1434B coupled with the compute clusters 1436A-1436H via a set of memory controllers 1442A-1442B. In various embodiments, the memory 1434A-1434B can include various types of memory devices including dynamic random access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory.

In one embodiment the compute clusters 1436A-1436H each include a set of graphics cores, such as the graphics core 1400 of FIG. 14A, which can include multiple types of

integer and floating point logic units that can perform computational operations at a range of precisions including suited for machine learning computations. For example and in one embodiment at least a subset of the floating point units in each of the compute clusters **1436A-1436H** can be configured to perform 16-bit or 32-bit floating point operations, while a different subset of the floating point units can be configured to perform 64-bit floating point operations.

Multiple instances of the GPGPU **1430** can be configured to operate as a compute cluster. The communication mechanism used by the compute cluster for synchronization and data exchange varies across embodiments. In one embodiment the multiple instances of the GPGPU **1430** communicate over the host interface **1432**. In one embodiment the GPGPU **1430** includes an I/O hub **1439** that couples the GPGPU **1430** with a GPU link **1440** that enables a direct connection to other instances of the GPGPU. In one embodiment the GPU link **1440** is coupled to a dedicated GPU-to-GPU bridge that enables communication and synchronization between multiple instances of the GPGPU **1430**. In one embodiment the GPU link **1440** couples with a high speed interconnect to transmit and receive data to other GPGPUs or parallel processors. In one embodiment the multiple instances of the GPGPU **1430** are located in separate data processing systems and communicate via a network device that is accessible via the host interface **1432**. In one embodiment the GPU link **1440** can be configured to enable a connection to a host processor in addition to or as an alternative to the host interface **1432**.

While the illustrated configuration of the GPGPU **1430** can be configured to train neural networks, one embodiment provides alternate configuration of the GPGPU **1430** that can be configured for deployment within a high performance or low power inferencing platform. In an inferencing configuration the GPGPU **1430** includes fewer of the compute clusters **1436A-1436H** relative to the training configuration. Additionally, the memory technology associated with the memory **1434A-1434B** may differ between inferencing and training configurations, with higher bandwidth memory technologies devoted to training configurations. In one embodiment the inferencing configuration of the GPGPU **1430** can support inferencing specific instructions. For example, an inferencing configuration can provide support for one or more 8-bit integer dot product instructions, which are commonly used during inferencing operations for deployed neural networks

#### Machine Learning Overview

A machine learning algorithm is an algorithm that can learn based on a set of data. Embodiments of machine learning algorithms can be designed to model high-level abstractions within a data set. For example, image recognition algorithms can be used to determine which of several categories to which a given input belong; regression algorithms can output a numerical value given an input; and pattern recognition algorithms can be used to generate translated text or perform text to speech and/or speech recognition.

An exemplary type of machine learning algorithm is a neural network. There are many types of neural networks; a simple type of neural network is a feedforward network. A feedforward network may be implemented as an acyclic graph in which the nodes are arranged in layers. Typically, a feedforward network topology includes an input layer and an output layer that are separated by at least one hidden layer. The hidden layer transforms input received by the input layer into a representation that is useful for generating output in the output layer. The network nodes are fully

connected via edges to the nodes in adjacent layers, but there are no edges between nodes within each layer. Data received at the nodes of an input layer of a feedforward network are propagated (i.e., “fed forward”) to the nodes of the output layer via an activation function that calculates the states of the nodes of each successive layer in the network based on coefficients (“weights”) respectively associated with each of the edges connecting the layers. Depending on the specific model being represented by the algorithm being executed, the output from the neural network algorithm can take various forms.

Before a machine learning algorithm can be used to model a particular problem, the algorithm is trained using a training data set. Training a neural network involves selecting a network topology, using a set of training data representing a problem being modeled by the network, and adjusting the weights until the network model performs with a minimal error for all instances of the training data set. For example, during a supervised learning training process for a neural network, the output produced by the network in response to the input representing an instance in a training data set is compared to the “correct” labeled output for that instance, an error signal representing the difference between the output and the labeled output is calculated, and the weights associated with the connections are adjusted to minimize that error as the error signal is backward propagated through the layers of the network. The network is considered “trained” when the errors for each of the outputs generated from the instances of the training data set are minimized.

The accuracy of a machine learning algorithm can be affected significantly by the quality of the data set used to train the algorithm. The training process can be computationally intensive and may require a significant amount of time on a conventional general-purpose processor. Accordingly, parallel processing hardware is used to train many types of machine learning algorithms. This is particularly useful for optimizing the training of neural networks, as the computations performed in adjusting the coefficients in neural networks lend themselves naturally to parallel implementations. Specifically, many machine learning algorithms and software applications have been adapted to make use of the parallel processing hardware within general-purpose graphics processing devices.

FIG. **15** is a generalized diagram of a machine learning software stack **1500**. A machine learning application **1502** can be configured to train a neural network using a training dataset or to use a trained deep neural network to implement machine intelligence. The machine learning application **1502** can include training and inference functionality for a neural network and/or specialized software that can be used to train a neural network before deployment. The machine learning application **1502** can implement any type of machine intelligence including but not limited to image recognition, mapping and localization, autonomous navigation, speech synthesis, medical imaging, or language translation.

Hardware acceleration for the machine learning application **1502** can be enabled via a machine learning framework **1504**. The machine learning framework **1504** can provide a library of machine learning primitives. Machine learning primitives are basic operations that are commonly performed by machine learning algorithms. Without the machine learning framework **1504**, developers of machine learning algorithms would be required to create and optimize the main computational logic associated with the machine learning algorithm, then re-optimize the computational logic as new parallel processors are developed.

Instead, the machine learning application can be configured to perform the necessary computations using the primitives provided by the machine learning framework **1504**. Exemplary primitives include tensor convolutions, activation functions, and pooling, which are computational operations that are performed while training a convolutional neural network (CNN). The machine learning framework **1504** can also provide primitives to implement basic linear algebra subprograms performed by many machine-learning algorithms, such as matrix and vector operations.

The machine learning framework **1504** can process input data received from the machine learning application **1502** and generate the appropriate input to a compute framework **1506**. The compute framework **1506** can abstract the underlying instructions provided to the GPGPU driver **1508** to enable the machine learning framework **1504** to take advantage of hardware acceleration via the GPGPU hardware **1510** without requiring the machine learning framework **1504** to have intimate knowledge of the architecture of the GPGPU hardware **1510**. Additionally, the compute framework **1506** can enable hardware acceleration for the machine learning framework **1504** across a variety of types and generations of the GPGPU hardware **1510**.

#### Machine Learning Neural Network Implementations

The computing architecture provided by embodiments described herein can be configured to perform the types of parallel processing that is particularly suited for training and deploying neural networks for machine learning. A neural network can be generalized as a network of functions having a graph relationship. As is known in the art, there are a variety of types of neural network implementations used in machine learning. One exemplary type of neural network is the feedforward network, as previously described.

A second exemplary type of neural network is the Convolutional Neural Network (CNN). A CNN is a specialized feedforward neural network for processing data having a known, grid-like topology, such as image data. Accordingly, CNNs are commonly used for compute vision and image recognition applications, but they also may be used for other types of pattern recognition such as speech and language processing. The nodes in the CNN input layer are organized into a set of “filters” (feature detectors inspired by the receptive fields found in the retina), and the output of each set of filters is propagated to nodes in successive layers of the network. The computations for a CNN include applying the convolution mathematical operation to each filter to produce the output of that filter. Convolution is a specialized kind of mathematical operation performed by two functions to produce a third function that is a modified version of one of the two original functions. In convolutional network terminology, the first function to the convolution can be referred to as the input, while the second function can be referred to as the convolution kernel. The output may be referred to as the feature map. For example, the input to a convolution layer can be a multidimensional array of data that defines the various color components of an input image. The convolution kernel can be a multidimensional array of parameters, where the parameters are adapted by the training process for the neural network.

Recurrent neural networks (RNNs) are a family of feedforward neural networks that include feedback connections between layers. RNNs enable modeling of sequential data by sharing parameter data across different parts of the neural network. The architecture for a RNN includes cycles. The cycles represent the influence of a present value of a variable on its own value at a future time, as at least a portion of the output data from the RNN is used as feedback for processing

subsequent input in a sequence. This feature makes RNNs particularly useful for language processing due to the variable nature in which language data can be composed.

The figures described below present exemplary feedforward, CNN, and RNN networks, as well as describe a general process for respectively training and deploying each of those types of networks. It will be understood that these descriptions are exemplary and non-limiting as to any specific embodiment described herein and the concepts illustrated can be applied generally to deep neural networks and machine learning techniques in general.

The exemplary neural networks described above can be used to perform deep learning. Deep learning is machine learning using deep neural networks. The deep neural networks used in deep learning are artificial neural networks composed of multiple hidden layers, as opposed to shallow neural networks that include only a single hidden layer. Deeper neural networks are generally more computationally intensive to train. However, the additional hidden layers of the network enable multistep pattern recognition that results in reduced output error relative to shallow machine learning techniques.

Deep neural networks used in deep learning typically include a front-end network to perform feature recognition coupled to a back-end network which represents a mathematical model that can perform operations (e.g., object classification, speech recognition, etc.) based on the feature representation provided to the model. Deep learning enables machine learning to be performed without requiring hand crafted feature engineering to be performed for the model. Instead, deep neural networks can learn features based on statistical structure or correlation within the input data. The learned features can be provided to a mathematical model that can map detected features to an output. The mathematical model used by the network is generally specialized for the specific task to be performed, and different models will be used to perform different task.

Once the neural network is structured, a learning model can be applied to the network to train the network to perform specific tasks. The learning model describes how to adjust the weights within the model to reduce the output error of the network. Backpropagation of errors is a common method used to train neural networks. An input vector is presented to the network for processing. The output of the network is compared to the desired output using a loss function and an error value is calculated for each of the neurons in the output layer. The error values are then propagated backwards until each neuron has an associated error value which roughly represents its contribution to the original output. The network can then learn from those errors using an algorithm, such as the stochastic gradient descent algorithm, to update the weights of the of the neural network.

FIG. **16A-16B** illustrate an exemplary convolutional neural network. FIG. **16A** illustrates various layers within a CNN. As shown in FIG. **16A**, an exemplary CNN used to model image processing can receive input **1602** describing the red, green, and blue (RGB) components of an input image. The input **1602** can be processed by multiple convolutional layers (e.g., first convolutional layer **1604**, second convolutional layer **1606**). The output from the multiple convolutional layers may optionally be processed by a set of fully connected layers **1608**. Neurons in a fully connected layer have full connections to all activations in the previous layer, as previously described for a feedforward network. The output from the fully connected layers **1608** can be used to generate an output result from the network. The activations within the fully connected layers **1608** can be com-

puted using matrix multiplication instead of convolution. Not all CNN implementations make use of fully connected layers **1608**. For example, in some implementations the second convolutional layer **1606** can generate output for the CNN.

The convolutional layers are sparsely connected, which differs from traditional neural network configuration found in the fully connected layers **1608**. Traditional neural network layers are fully connected, such that every output unit interacts with every input unit. However, the convolutional layers are sparsely connected because the output of the convolution of a field is input (instead of the respective state value of each of the nodes in the field) to the nodes of the subsequent layer, as illustrated. The kernels associated with the convolutional layers perform convolution operations, the output of which is sent to the next layer. The dimensionality reduction performed within the convolutional layers is one aspect that enables the CNN to scale to process large images.

FIG. **16B** illustrates exemplary computation stages within a convolutional layer of a CNN. Input to a convolutional layer **1612** of a CNN can be processed in three stages of a convolutional layer **1614**. The three stages can include a convolution stage **1616**, a detector stage **1618**, and a pooling stage **1620**. The convolution layer **1614** can then output data to a successive convolutional layer. The final convolutional layer of the network can generate output feature map data or provide input to a fully connected layer, for example, to generate a classification value for the input to the CNN.

In the convolution stage **1616** performs several convolutions in parallel to produce a set of linear activations. The convolution stage **1616** can include an affine transformation, which is any transformation that can be specified as a linear transformation plus a translation. Affine transformations include rotations, translations, scaling, and combinations of these transformations. The convolution stage computes the output of functions (e.g., neurons) that are connected to specific regions in the input, which can be determined as the local region associated with the neuron. The neurons compute a dot product between the weights of the neurons and the region in the local input to which the neurons are connected. The output from the convolution stage **1616** defines a set of linear activations that are processed by successive stages of the convolutional layer **1614**.

The linear activations can be processed by a detector stage **1618**. In the detector stage **1618**, each linear activation is processed by a non-linear activation function. The non-linear activation function increases the nonlinear properties of the overall network without affecting the receptive fields of the convolution layer. Several types of non-linear activation functions may be used. One particular type is the rectified linear unit (ReLU), which uses an activation function defined as  $f(x)=\max(0,x)$ , such that the activation is thresholded at zero.

The pooling stage **1620** uses a pooling function that replaces the output of the second convolutional layer **1606** with a summary statistic of the nearby outputs. The pooling function can be used to introduce translation invariance into the neural network, such that small translations to the input do not change the pooled outputs. Invariance to local translation can be useful in scenarios where the presence of a feature in the input data is more important than the precise location of the feature. Various types of pooling functions can be used during the pooling stage **1620**, including max pooling, average pooling, and 12-norm pooling. Additionally, some CNN implementations do not include a pooling stage. Instead, such implementations substitute and addi-

tional convolution stage having an increased stride relative to previous convolution stages.

The output from the convolutional layer **1614** can then be processed by the next layer **1622**. The next layer **1622** can be an additional convolutional layer or one of the fully connected layers **1608**. For example, the first convolutional layer **1604** of FIG. **16A** can output to the second convolutional layer **1606**, while the second convolutional layer can output to a first layer of the fully connected layers **1608**.

FIG. **17** illustrates an exemplary recurrent neural network. In a recurrent neural network (RNN), the previous state of the network influences the output of the current state of the network. RNNs can be built in a variety of ways using a variety of functions. The use of RNNs generally revolves around using mathematical models to predict the future based on a prior sequence of inputs. For example, an RNN may be used to perform statistical language modeling to predict an upcoming word given a previous sequence of words. The illustrated RNN **1700** can be described as having an input layer **1702** that receives an input vector, hidden layers **1704** to implement a recurrent function, a feedback mechanism **1705** to enable a 'memory' of previous states, and an output layer **1706** to output a result. The RNN **1700** operates based on time-steps. The state of the RNN at a given time step is influenced based on the previous time step via the feedback mechanism **1705**. For a given time step, the state of the hidden layers **1704** is defined by the previous state and the input at the current time step. An initial input ( $x_1$ ) at a first time step can be processed by the hidden layer **1704**. A second input ( $x_2$ ) can be processed by the hidden layer **1704** using state information that is determined during the processing of the initial input ( $x_1$ ). A given state can be computed as  $s_t=f(Ux_t+Ws_{t-1})$ , where  $U$  and  $W$  are parameter matrices. The function  $f$  is generally a nonlinearity, such as the hyperbolic tangent function (Tanh) or a variant of the rectifier function  $f(x)=\max(0,x)$ . However, the specific mathematical function used in the hidden layers **1704** can vary depending on the specific implementation details of the RNN **1700**.

In addition to the basic CNN and RNN networks described, variations on those networks may be enabled. One example RNN variant is the long short-term memory (LSTM) RNN. LSTM RNNs are capable of learning long-term dependencies that may be necessary for processing longer sequences of language. A variant on the CNN is a convolutional deep belief network, which has a structure similar to a CNN and is trained in a manner similar to a deep belief network. A deep belief network (DBN) is a generative neural network that is composed of multiple layers of stochastic (random) variables. DBNs can be trained layer-by-layer using greedy unsupervised learning. The learned weights of the DBN can then be used to provide pre-train neural networks by determining an optimal initial set of weights for the neural network.

FIG. **18** illustrates training and deployment of a deep neural network. Once a given network has been structured for a task the neural network is trained using a training dataset **1802**. Various training frameworks have been developed to enable hardware acceleration of the training process. For example, the machine learning framework **1504** of FIG. **15** may be configured as a training framework **1804**. The training framework **1804** can hook into an untrained neural network **1806** and enable the untrained neural net to be trained using the parallel processing resources described herein to generate a trained neural network **1808**. To start the training process the initial weights may be chosen randomly

or by pre-training using a deep belief network. The training cycle then be performed in either a supervised or unsupervised manner.

Supervised learning is a learning method in which training is performed as a mediated operation, such as when the training dataset **1802** includes input paired with the desired output for the input, or where the training dataset includes input having known output and the output of the neural network is manually graded. The network processes the inputs and compares the resulting outputs against a set of expected or desired outputs. Errors are then propagated back through the system. The training framework **1804** can adjust to adjust the weights that control the untrained neural network **1806**. The training framework **1804** can provide tools to monitor how well the untrained neural network **1806** is converging towards a model suitable to generating correct answers based on known input data. The training process occurs repeatedly as the weights of the network are adjusted to refine the output generated by the neural network. The training process can continue until the neural network reaches a statistically desired accuracy associated with a trained neural network **1808**. The trained neural network **1808** can then be deployed to implement any number of machine learning operations.

Unsupervised learning is a learning method in which the network attempts to train itself using unlabeled data. Thus, for unsupervised learning the training dataset **1802** will include input data without any associated output data. The untrained neural network **1806** can learn groupings within the unlabeled input and can determine how individual inputs are related to the overall dataset. Unsupervised training can be used to generate a self-organizing map, which is a type of trained neural network **1807** capable of performing operations useful in reducing the dimensionality of data. Unsupervised training can also be used to perform anomaly detection, which allows the identification of data points in an input dataset that deviate from the normal patterns of the data.

Variations on supervised and unsupervised training may also be employed. Semi-supervised learning is a technique in which in the training dataset **1802** includes a mix of labeled and unlabeled data of the same distribution. Incremental learning is a variant of supervised learning in which input data is continuously used to further train the model. Incremental learning enables the trained neural network **1808** to adapt to the new data **1812** without forgetting the knowledge instilled within the network during initial training.

Whether supervised or unsupervised, the training process for particularly deep neural networks may be too computationally intensive for a single compute node. Instead of using a single compute node, a distributed network of computational nodes can be used to accelerate the training process.

FIG. **19** is a block diagram illustrating distributed learning. Distributed learning is a training model that uses multiple distributed computing nodes to perform supervised or unsupervised training of a neural network. The distributed computational nodes can each include one or more host processors and one or more of the general-purpose processing nodes. As illustrated, distributed learning can be performed model parallelism **1902**, data parallelism **1904**, or a combination of model and data parallelism **1904**.

In model parallelism **1902**, different computational nodes in a distributed system can perform training computations for different parts of a single network. For example, each layer of a neural network can be trained by a different processing node of the distributed system. The benefits of

model parallelism include the ability to scale to particularly large models. Splitting the computations associated with different layers of the neural network enables the training of very large neural networks in which the weights of all layers would not fit into the memory of a single computational node. In some instances, model parallelism can be particularly useful in performing unsupervised training of large neural networks.

In data parallelism **1904**, the different nodes of the distributed network have a complete instance of the model and each node receives a different portion of the data. The results from the different nodes are then combined. While different approaches to data parallelism are possible, data parallel training approaches all require a technique of combining results and synchronizing the model parameters between each node. Exemplary approaches to combining data include parameter averaging and update based data parallelism. Parameter averaging trains each node on a subset of the training data and sets the global parameters (e.g., weights, biases) to the average of the parameters from each node. Parameter averaging uses a central parameter server that maintains the parameter data. Update based data parallelism is similar to parameter averaging except that instead of transferring parameters from the nodes to the parameter server, the updates to the model are transferred. Additionally, update based data parallelism can be performed in a decentralized manner, where the updates are compressed and transferred between nodes.

Combined model and data parallelism **1906** can be implemented, for example, in a distributed system in which each computational node includes multiple GPUs. Each node can have a complete instance of the model with separate GPUs within each node are used to train different portions of the model.

Distributed training has increased overhead relative to training on a single machine. However, the parallel processors and GPGPUs described herein can each implement various techniques to reduce the overhead of distributed training, including techniques to enable high bandwidth GPU-to-GPU data transfer and accelerated remote data synchronization.

#### Exemplary Machine Learning Applications

Machine learning can be applied to solve a variety of technological problems, including but not limited to computer vision, autonomous driving and navigation, speech recognition, and language processing. Computer vision has traditionally been one of the most active research areas for machine learning applications. Applications of computer vision range from reproducing human visual abilities, such as recognizing faces, to creating new categories of visual abilities. For example, computer vision applications can be configured to recognize sound waves from the vibrations induced in objects visible in a video. Parallel processor accelerated machine learning enables computer vision applications to be trained using significantly larger training dataset than previously feasible and enables inferencing systems to be deployed using low power parallel processors.

Parallel processor accelerated machine learning has autonomous driving applications including lane and road sign recognition, obstacle avoidance, navigation, and driving control. Accelerated machine learning techniques can be used to train driving models based on datasets that define the appropriate responses to specific training input. The parallel processors described herein can enable rapid training of the increasingly complex neural networks used for autonomous driving solutions and enables the deployment of low power

inferencing processors in a mobile platform suitable for integration into autonomous vehicles.

Parallel processor accelerated deep neural networks have enabled machine learning approaches to automatic speech recognition (ASR). ASR includes the creation of a function that computes the most probable linguistic sequence given an input acoustic sequence. Accelerated machine learning using deep neural networks have enabled the replacement of the hidden Markov models (HMMs) and Gaussian mixture models (GMMs) previously used for ASR.

Parallel processor accelerated machine learning can also be used to accelerate natural language processing. Automatic learning procedures can make use of statistical inference algorithms to produce models that are robust to erroneous or unfamiliar input. Exemplary natural language processor applications include automatic machine translation between human languages.

The parallel processing platforms used for machine learning can be divided into training platforms and deployment platforms. Training platforms are generally highly parallel and include optimizations to accelerate multi-GPU single node training and multi-node, multi-GPU training, while deployed machine learning (e.g., inferencing) platforms generally include lower power parallel processors suitable for use in products such as cameras, autonomous robots, and autonomous vehicles.

FIG. 20 illustrates one embodiment of a computing device 2000 employing an accelerator mechanism. Computing device 2000 (e.g., smart wearable devices, virtual reality (VR) devices, head-mounted display (HMDs), mobile computers, Internet of Things (IoT) devices, laptop computers, desktop computers, server computers, etc.) may be the same as data processing system 100 of FIG. 1 and accordingly, for brevity, clarity, and ease of understanding, many of the details stated above with reference to FIGS. 1-19 are not further discussed or repeated hereafter. As illustrated, in one embodiment, computing device 2000 is shown as hosting an accelerator mechanism 2010.

Although illustrated as a separate component, other embodiments may feature accelerator mechanism 2010 being hosted by graphics processing unit (GPU) 2014. In other embodiments, accelerator mechanism 2010 may be hosted by or part of firmware of graphics driver 2016. In yet other embodiments, accelerator mechanism 2010 may be hosted by or part of firmware of central processing unit (“CPU” or “application processor”) 2012. For brevity, clarity, and ease of understanding, throughout the rest of this document, accelerator mechanism 2010 may be discussed as part of GPU 2014; however, embodiments are not limited as such.

In yet another embodiment, accelerator mechanism 2010 may be hosted as software or firmware logic by operating system 2006. In yet a further embodiment, accelerator mechanism 2010 may be partially and simultaneously hosted by multiple components of computing device 2000, such as one or more of graphics driver 616, GPU 2014, GPU firmware, CPU 2012, CPU firmware, operating system 2006, and/or the like. It is contemplated that accelerator mechanism 2010 or one or more of their components may be implemented as hardware, software, and/or firmware.

Throughout the document, term “user” may be interchangeably referred to as “viewer”, “observer”, “person”, “individual”, “end-user”, and/or the like. It is to be noted that throughout this document, terms like “graphics domain” may be referenced interchangeably with “graphics processing unit”, “graphics processor”, or simply “GPU” and similarly, “CPU domain” or “host domain” may be referenced

interchangeably with “computer processing unit”, “application processor”, or simply “CPU”. Computing device 2000 may include any number and type of communication devices, such as large computing systems, such as server computers, desktop computers, etc., and may further include set-top boxes (e.g., Internet-based cable television set-top boxes, etc.), global positioning system (GPS)-based devices, etc. Computing device 2000 may include mobile computing devices serving as communication devices, such as cellular phones including smartphones, personal digital assistants (PDAs), tablet computers, laptop computers, e-readers, smart televisions, television platforms, wearable devices (e.g., glasses, watches, bracelets, smartcards, jewelry, clothing items, etc.), media players, etc. For example, in one embodiment, computing device 2000 may include a mobile computing device employing a computer platform hosting an integrated circuit (“IC”), such as system on a chip (“SoC” or “SOC”), integrating various hardware and/or software components of computing device 2000 on a single chip.

As illustrated, in one embodiment, computing device 2000 may include any number and type of hardware and/or software components, such as (without limitation) GPU 2014, graphics driver (also referred to as “GPU driver”, “graphics driver logic”, “driver logic”, user-mode driver (UMD), UMD, user-mode driver framework (UMDF), UMDF, or simply “driver”) 616, CPU 2012, memory 2008, network devices, drivers, or the like, as well as input/output (I/O) sources 2004, such as touchscreens, touch panels, touch pads, virtual or regular keyboards, virtual or regular mice, ports, connectors, etc.

Computing device 2000 may include operating system (OS) 2006 serving as an interface between hardware and/or physical resources of the computer device 2000 and a user. It is contemplated that CPU 2012 may include one or more processors, such as processor(s) 102 of FIG. 1, while GPU 2014 may include one or more graphics processors (or multiprocessors).

It is to be noted that terms like “node”, “computing node”, “server”, “server device”, “cloud computer”, “cloud server”, “cloud server computer”, “machine”, “host machine”, “device”, “computing device”, “computer”, “computing system”, and the like, may be used interchangeably throughout this document. It is to be further noted that terms like “application”, “software application”, “program”, “software program”, “package”, “software package”, and the like, may be used interchangeably throughout this document. Also, terms like “job”, “input”, “request”, “message”, and the like, may be used interchangeably throughout this document.

It is contemplated and as further described with reference to FIGS. 1-14, some processes of the graphics pipeline as described above are implemented in software, while the rest are implemented in hardware. A graphics pipeline may be implemented in a graphics coprocessor design, where CPU 2012 is designed to work with GPU 2014 which may be included in or co-located with CPU 2012. In one embodiment, GPU 2014 may employ any number and type of conventional software and hardware logic to perform the conventional functions relating to graphics rendering as well as novel software and hardware logic to execute any number and type of instructions.

As aforementioned, memory 2008 may include a random access memory (RAM) comprising application database having object information. A memory controller hub, such as memory hub 105 of FIG. 1, may access data in the RAM and forward it to GPU 2014 for graphics pipeline processing. RAM may include double data rate RAM (DDR RAM),

extended data output RAM (EDO RAM), etc. CPU **2012** interacts with a hardware graphics pipeline to share graphics pipelining functionality.

Processed data is stored in a buffer in the hardware graphics pipeline, and state information is stored in memory **2008**. The resulting image is then transferred to I/O sources **2004**, such as a display component for displaying of the image. It is contemplated that the display device may be of various types, such as Cathode Ray Tube (CRT), Thin Film Transistor (TFT), Liquid Crystal Display (LCD), Organic Light Emitting Diode (OLED) array, etc., to display information to a user.

Memory **2008** may comprise a pre-allocated region of a buffer (e.g., frame buffer); however, it should be understood by one of ordinary skill in the art that the embodiments are not so limited, and that any memory accessible to the lower graphics pipeline may be used. Computing device **2000** may further include input/output (I/O) control hub (ICH) **107** as referenced in FIG. 1, as one or more I/O sources **2004**, etc.

CPU **2012** may include one or more processors to execute instructions in order to perform whatever software routines the computing system implements. The instructions frequently involve some sort of operation performed upon data. Both data and instructions may be stored in system memory **2008** and any associated cache. Cache is typically designed to have shorter latency times than system memory **2008**; for example, cache might be integrated onto the same silicon chip(s) as the processor(s) and/or constructed with faster static RAM (SRAM) cells whilst the system memory **2008** might be constructed with slower dynamic RAM (DRAM) cells. By tending to store more frequently used instructions and data in the cache as opposed to the system memory **2008**, the overall performance efficiency of computing device **2000** improves. It is contemplated that in some embodiments, GPU **2014** may exist as part of CPU **2012** (such as part of a physical CPU package) in which case, memory **2008** may be shared by CPU **2012** and GPU **2014** or kept separated.

System memory **2008** may be made available to other components within the computing device **2000**. For example, any data (e.g., input graphics data) received from various interfaces to the computing device **2000** (e.g., keyboard and mouse, printer port, Local Area Network (LAN) port, modem port, etc.) or retrieved from an internal storage element of the computer device **2000** (e.g., hard disk drive) are often temporarily queued into system memory **2008** prior to being operated upon by the one or more processor(s) in the implementation of a software program. Similarly, data that a software program determines should be sent from the computing device **2000** to an outside entity through one of the computing system interfaces, or stored into an internal storage element, is often temporarily queued in system memory **2008** prior to its being transmitted or stored.

Further, for example, an ICH may be used for ensuring that such data is properly passed between the system memory **2008** and its appropriate corresponding computing system interface (and internal storage device if the computing system is so designed) and may have bi-directional point-to-point links between itself and the observed **110** sources/devices **2004**. Similarly, an MCH may be used for managing the various contending requests for system memory **2008** accesses amongst CPU **2012** and GPU **2014**, interfaces and internal storage elements that may proximately arise in time with respect to one another.

I/O sources **2004** may include one or more I/O devices that are implemented for transferring data to and/or from computing device **2000** (e.g., a networking adapter); or, for

a large scale non-volatile storage within computing device **2000** (e.g., hard disk drive). User input device, including alphanumeric and other keys, may be used to communicate information and command selections to GPU **2014**. Another type of user input device is cursor control, such as a mouse, a trackball, a touchscreen, a touchpad, or cursor direction keys to communicate direction information and command selections to GPU **2014** and to control cursor movement on the display device. Camera and microphone arrays of computer device **2000** may be employed to observe gestures, record audio and video and to receive and transmit visual and audio commands.

Computing device **2000** may further include network interface(s) to provide access to a network, such as a LAN, a wide area network (WAN), a metropolitan area network (MAN), a personal area network (PAN), Bluetooth, a cloud network, a mobile network (e.g., 3rd Generation (3G), 4th Generation (4G), etc.), an intranet, the Internet, etc. Network interface(s) may include, for example, a wireless network interface having antenna, which may represent one or more antenna(e). Network interface(s) may also include, for example, a wired network interface to communicate with remote devices via network cable, which may be, for example, an Ethernet cable, a coaxial cable, a fiber optic cable, a serial cable, or a parallel cable.

Network interface(s) may provide access to a LAN, for example, by conforming to IEEE 802.11b and/or IEEE 802.11g standards, and/or the wireless network interface may provide access to a personal area network, for example, by conforming to Bluetooth standards. Other wireless network interfaces and/or protocols, including previous and subsequent versions of the standards, may also be supported. In addition to, or instead of, communication via the wireless LAN standards, network interface(s) may provide wireless communication using, for example, Time Division, Multiple Access (TDMA) protocols, Global Systems for Mobile Communications (GSM) protocols, Code Division, Multiple Access (CDMA) protocols, and/or any other type of wireless communications protocols.

Network interface(s) may include one or more communication interfaces, such as a modem, a network interface card, or other well-known interface devices, such as those used for coupling to the Ethernet, token ring, or other types of physical wired or wireless attachments for purposes of providing a communication link to support a LAN or a WAN, for example. In this manner, the computer system may also be coupled to a number of peripheral devices, clients, control surfaces, consoles, or servers via a conventional network infrastructure, including an Intranet or the Internet, for example.

It is to be appreciated that a lesser or more equipped system than the example described above may be preferred for certain implementations. Therefore, the configuration of computing device **2000** may vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, technological improvements, or other circumstances. Examples of the electronic device or computer system **2000** may include (without limitation) a mobile device, a personal digital assistant, a mobile computing device, a smartphone, a cellular telephone, a handset, a one-way pager, a two-way pager, a messaging device, a computer, a personal computer (PC), a desktop computer, a laptop computer, a notebook computer, a handheld computer, a tablet computer, a server, a server array or server farm, a web server, a network server, an Internet server, a work station, a mini-computer, a main frame computer, a supercomputer, a network appliance, a

web appliance, a distributed computing system, multiprocessor systems, processor-based systems, consumer electronics, programmable consumer electronics, television, digital television, set top box, wireless access point, base station, subscriber station, mobile subscriber center, radio network controller, router, hub, gateway, bridge, switch, machine, or combinations thereof.

Embodiments may be implemented as any or a combination of: one or more microchips or integrated circuits interconnected using a parentboard, hardwired logic, software stored by a memory device and executed by a microprocessor, firmware, an application specific integrated circuit (ASIC), and/or a field programmable gate array (FPGA). The term "logic" may include, by way of example, software or hardware and/or combinations of software and hardware.

Embodiments may be provided, for example, as a computer program product which may include one or more machine-readable media having stored thereon machine-executable instructions that, when executed by one or more machines such as a computer, network of computers, or other electronic devices, may result in the one or more machines carrying out operations in accordance with embodiments described herein. A machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs (Compact Disc-Read Only Memories), and magneto-optical disks, ROMs, RAMs, EPROMs (Erasable Programmable Read Only Memories), EEPROMs (Electrically Erasable Programmable Read Only Memories), magnetic or optical cards, flash memory, or other type of media/machine-readable medium suitable for storing machine-executable instructions.

Moreover, embodiments may be downloaded as a computer program product, wherein the program may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of one or more data signals embodied in and/or modulated by a carrier wave or other propagation medium via a communication link (e.g., a modem and/or network connection).

Performing machine learning (e.g. deep learning and neural networks) applications at computing device **2000**. Specifically, such applications are typically processed at a GPU, such as GPU **2014**. However, the compute intensive nature of machine learning applications may result in decreased GPU performance. For instance, machine learning algorithms typically include a graph structure in which data flows from one node to the next. Each node represents an operation (or op) executed on the data structure provided to the graph. Various types of parallelization are exploited in the graph, including: 1) parallelization inside the node; 2) parallel execution of multiple nodes; 3) forward or backward pass for a convolutional neural network (CNN); and 4) data communication for a distributed algorithm. The hardware resources (e.g., memory channel, cache, number of cores, frequencies, threads, power, level of accuracy (or precision), etc.) needed for the parallel execution of the graph nodes varies based on the above cases. For example, hundreds of thousands of threads are available in a GPU, and if not properly scheduled, will be under-subscribed or over-subscribed, which leads to poor performance of running the machine learning algorithm.

According to various embodiments, an acceleration mechanism **2010** is implemented to improve computing device **2000** performance for machine learning applications. In one such an embodiment, a fast interconnect acceleration mechanism **2010** may incorporate a fast interconnect between a GPU and a CPU (e.g., CPU **2012**) to accelerate

deep learning tasks. In this embodiment, acceleration mechanism **2010** is implemented in a Differentiable Neural Computer (DNC). A DNC is a recurrent artificial neural network architecture with an auto associative memory. Thus, a DNC operates as a neural network with external memory that stores knowledge as opposed to state. FIG. **21A** illustrates an example DNC **2100**.

As shown in FIG. **21A**, DNC **2100** includes a memory **2101**, external controller **2102** and input/output (I/O) **2103**. As discussed above, memory **2101** is an external memory used to store knowledge data. Controller **2102** is a neural network implemented by Long Short-Term Memory (LSTM). Controller **2102** is responsible for reading input in, reading from and writing to memory, and producing output that may be interpreted as an answer. Memory **2101** is a set of locations that can each store a vector of information.

Controller **2102** can perform several operations on memory **2101** to update what is stored at a location. When the controller **2102** writes, it sends a vector of information to the chosen location in memory. Each time data is written, the locations are connected by links of association, which represent the order in which information was stored. Controller **2102** may also read from multiple locations in memory **2101**. Memory **2101** can be searched based on the content of each location, or the associative temporal links can be followed forward and backward to recall information written in sequence or in reverse. The read out information can be used to produce answers to questions or actions to take in an environment.

In one embodiment, external I/O **2102** transforms raw input (e.g., image, sentence, signal, action) to a vector of bits and transforms a vector of bits back to raw output, as shown in FIG. **21A**. I/O **2102** may implement different methods of transformation, such as branching capabilities, and general compute capabilities that are difficult to program on controller **2103**. According to one embodiment, external controller **2102** is implemented by GPU **2014**, while CPU **2012** is implemented to perform the transformations. In a further embodiment, memory **2101** is implemented via a shared memory between CPU **2012** and GPU **2014**, which uses a GPU zero-copy feature to facilitate communication between CPU **2012** and GPU **2014**. Thus, shared memory **2101** allows for a quick transfer to receive data from the GPU **2014** neural network to external memory.

FIG. **21B** illustrates one embodiment of a DNC flow in which CPU **2012** provides input to the GPU **2014** neural network, which then forwards output back to CPU **2012**. Mapping each of the DNC sub-components to either CPU **2012** or GPU **2014** per their relative strengths, while sharing the same memory with minimal overhead, leads to an optimal implementation. The above-described embodiment cannot currently be implemented on standard deep learning GPU only systems due to the reliance on external memory given the need for general compute capabilities required by some DNC components.

In a further embodiment, acceleration mechanism **2010** includes an intelligent scheduler **2015** that statically analyzes a graph of a neural network algorithm and the hardware resources (e.g., multi CPU+GPU nodes) at computing device **2000**, and produces a map for assigning resources to process the graph nodes. Thus, scheduler **2015** enables faster model training and inference on multi CPU+GPU hardware.

In one embodiment, CPU **2012** and GPU **2014** operate individually serially or simultaneously in parallel. As discussed above, deep learning algorithms are compute intensive, and include algorithms, such as 2D Convolution, Rectifier Linear Unit (RELU), Batch Normalization, Matrix

Multiplications, and other operators that are optimized by various software and hardware vendors. The optimized operators are bundled in libraries such as Math Kernel Library Deep Neural Network (MKL DNN) by Intel Corporation®, (Compute Unified Deep Neural Network) cuDNN by NVIDIA®, and several other libraries by various software and hardware companies.

The above mentioned optimized libraries may be used by most Deep Learning (DL) frameworks, such as Tensorflow, MxNet, CNTK, etc. Most frameworks use a compute graph at their core, which may be described in deep learning may be described as nodes and edges. Nodes are deep learning operators, while edges represent a flow of data structures from one node to the next. For example, in Tensorflow, a simple graph includes nodes as 2D Convolutions and RELU, and the edge between 2D Convolution and RELU represents the flow of output tensors from the output of 2D Convolution to the input of RELU.

Although the Deep Neural Network (DNN) operators are optimized by hardware and software vendors, the performance may vary based on various types of CPUs and GPUs based on their hardware capabilities. For example, 2D Convolution may perform better on a GPU, while Batch Normalization can perform better on CPU. Accordingly, intelligent scheduler **2015** efficiently schedules operators based on the knowledge of a DNN compute graph, and the hardware capabilities of CPU and GPU.

In one embodiment, the hardware capabilities include, but are not limited to, compute throughput (as measured Giga Flops per second), memory bandwidth (Gigabytes per second), latency (Nanosecond), and several other hardware metrics. If the DNN operator is implemented in such a way that it can run parallel on CPU and GPU by using data parallelization, intelligent scheduler **2015** may schedule the operator to run on the CPU and GPU in parallel. Moreover, multiple subgraphs, which are independent and can run in parallel, may be included in the compute graph. Thus, intelligent scheduler **2015** efficiently schedules each of the subgraphs on the CPUs and GPUs available in the server.

FIG. **22** illustrates one embodiment of a scheduling process implemented at a machine learning acceleration mechanism. At processing block **2205**, computation costs of graph nodes and sub-graphs are determined. In one embodiment, scheduler **2015** performs this determination based on available information regarding a computation cost of the each of the DNN operators on CPU and GPU individually, and in parallel on CPU and GPU, if the operator is implemented to run simultaneously on both CPU and GPU. However, the computation cost of the each of the operators is again based on the hardware performance metrics of the CPU and GPU, and, the optimized complete compute graph.

At processing block **2206**, a least computation cost path is determined. In one embodiment, scheduler **2015** implements a shortest path algorithm (e.g., a Dijkstra shortest path algorithm) to the graph using the above-described computation cost metric to find the path with the least computation cost. In one embodiment, the least computation cost path is the map between the compute resource (CPU and GPU) and the compute graph. FIG. **23** illustrates one embodiment of a directed DNN model in representation map and a shortest path search implemented at a machine learning acceleration mechanism. As shown in FIG. **23**, the shortest path search determines the best path for a map including nodes **1-7** (e.g., whether an operator is mapped to the CPU, GPU or hybrid as both the CPU and the GPU). Moreover, FIG. **23** illustrates that the algorithm includes a search for each node as to whether CPU, GPU and hybrid

mapping provides the shortest path for the respective node. Thus, nodes **1-7** each include a path based on CPU, GPU and hybrid mapping.

At processing block **2207**, scheduler **2015** schedules the compute graph to the compute resource. In a further embodiment, a user may manually schedule the compute resource for the compute graph. In such an embodiment, the user manual scheduling takes precedence over the scheduler **2015** mapping. FIG. **24** illustrates one embodiment of a map scheduled by scheduler **2015**. As shown in FIG. **24**, nodes **1**, **3** and **4** include a C, denoting that the nodes are mapped to the CPU, while nodes **2**, **5** and **6** include a G, denoting that the nodes are mapped to the GPU. Node **7** includes an H, denoting that the node is hybrid mapped to the both the CPU and GPU.

The intelligent scheduler efficiently utilizes the available CPU and GPU hardware resources at a computing device to achieve the maximum performance of graph structure DNN models. Such a process is nearly impossible to achieve the same capability if a user wants to do it manually. Since to achieve that, the user needs to create and examine the graph, investigate the performance of each of the DNN operators of the graph on CPU and GPU, run the shortest path algorithm, and manually map the operators to the CPU and GPU before running the model. This manual process is time consuming and error prone. Thus, the intelligent scheduler can solve this scheduling problem efficiently.

Prior to implementation in a deep learning application, a neural network is typically trained. FIG. **25** illustrates one embodiment of deep learning topology **2511** including layers **2512(N)** and **2512(N+1)**. As shown in FIG. **25**, each layer **2512** includes a forward pass and a back-propagation pass. During the forward pass, activations are computed for each layer. The forward pass compute multiplies an input matrix with a weight (or model) matrix to achieve an output (or activation) matrix, which is provided as an input to a subsequent layer, such as layer **2512(N+1)**. FIG. **26** illustrates one embodiment of matrices and performed compute operations.

Once forward propagation is completed at all of the layers, back-propagation is performed from the last layer through to the first layer (e.g., layer **2512**). During the back-propagation process, back-propagation and weight gradient computes are performed, as shown in FIG. **25**. As illustrated above, training compute time is dominated by the forward propagation, propagation and weight gradient computes. Typically, the weight gradient compute is performed after the back-propagation compute due to limited resources available at a GPU. Moreover, when it is time to perform the weight gradient computes, the input data used during forward propagation must be retrieved from memory due to limited cache available at the CPU, further slowing the process.

According to one embodiment, accelerator mechanism **2010** implements a programmable accelerator to perform data layout transformations, such as the computation of weight gradient computes. In such an embodiment, weight gradient computes may be performed in parallel with the back-propagation computes performed at the GPU. FIG. **27** illustrates one embodiment of an accelerator **2708** integrated circuit (IC) including a memory **2710**. In one embodiment memory **2710** is a software controlled memory (e.g., cache or scratch pad) that receives input data for each layer used during the forward propagation compute.

As shown in FIG. **26**, a weight gradient compute multiplies a transpose of an input matrix with a weight matrix. In a further embodiment, accelerator **2708** computes the trans-

pose for each input matrix upon receiving the matrices during the forward propagation compute. Thus, the transpose of the input matrices are immediately available for compute upon receiving output matrices during back-propagation. In some embodiments, accelerator **2708** may be implemented on the same integrated circuit as the GPU. However other embodiments may feature accelerator **2708** on a separate integrated circuit.

In another embodiment, accelerator **2708** may be implemented to perform additional data deep layout transformations other than a transpose. For instance, data layout is one of the primary sources of inefficiencies in deep learning applications. Different architectures have different layout requirements forcing programmers to make specific choices. The idea here is to do these transformations on the fly and potentially do any data transformations required from one layer to another. Thus, data layout conversions (e.g., NHWC-NCHW data format conversions) may be offloaded to accelerator **2709** to enable parallel execution with other computations performed at the GPU.

In yet another embodiment, accelerator **2708** may be implemented to perform switching between layer-specific arithmetic precision (e.g. FP32 vs. FP16). In still another embodiment, accelerator **2709** may be implemented to perform data speculation to accelerate the above-described transformations and re-execute part or whole of the transformation based on partial availability of data. In the above-described embodiments, accelerator **2708** may be programmed at the start of execution to specify the transformations required from one layer to the next, with memory **2710** being used to hold intermediate data.

According to a further embodiment, accelerator **2708** may comprise a FPGA-ASIC (FA) integrated circuit architecture. In such an embodiment, the FA includes a heterogeneous multi-chip integration of FPGAs with domain-specific ASIC accelerators to enable the advantage of extreme flexibility of FPGAs as well as the top efficiency of custom-ASICs. In a further embodiment, the FPGAs and ASICs are integrated using Chip-on-Wafer-on-Substrate (CoWoS®) technology to integrate multiple chips side-by-side onto a interposer containing through-silicon vias (TSVs) by chip-to-wafer bonding process, which is followed by CoW chip assembly onto a substrate (CoW-on-Substrate). In still another embodiment, the FA may include multi-chip integration of GPUs domain-specific ASIC accelerators.

FIG. **28A** illustrates one embodiment of an FA accelerator **2809**, including features similar to accelerator **2708**, including an FPGA/GPU **2821**, ASICs **2822** and interfaces **2823** coupled between FPGA/GPU **2821** and ASICs **2822**. In one embodiment, interfaces **2823** are implemented via an Embedded Multi-die Interconnect Bridge (EMIB). An application domain typically includes one or more “mature” key operations that are shared across applications within the domain, and “immature” operations specific to each applications. In FIG. **28A** the mature shared key operations in the target application domain are implemented on custom ASIC accelerator integrated circuits **2822**, while the remaining application-specific (e.g., immature portion is implemented on FPGA **2821**. Thus, the overall integrated FA delivers significantly higher performance and efficiency than an “FPGA only” solution, while being extremely more flexible and programmable than an “ASIC only” solution. Furthermore, the FA accelerator **2809** allows for separation of concerns in the design, development, and choices of the ASIC, FPGA and/or GPU chips to use. Hence, no FPGA fabric change is required, which allows leveraging existing FPGA ecosystems (e.g., integrated circuits, platforms, soft-

ware, tools); and provides freedom in the ASIC design since it does not need to abide to FPGA/GPU fabric integration constraints.

Other FA configurations may be implemented based on the goals and requirements for a particular project (e.g., what domain being targeted, performance/energy requirements, process technology target(s), bandwidth/latency characteristics of multi-chip interfaces, etc. FIGS. **28B-28E** illustrate embodiments of FA accelerator **2809** in various other embodiments.

The FA accelerator **2809** interfaces to the FPGA/GPU fabric in the same manner as any other existing external hard FPGA/GPU IPs (e.g., PCIe). The ASIC connects to certain I/O pins on the FPGA/GPU. The FPGA/GPU sends commands and data via these pins to tell the ASIC what to do. The ASIC performs the requested commands and provides results back to FPGA/GPU. Hence, an FPGA/GPU design that interfaces with the ASIC exposes top-level external interface to connect to the appropriate FPGA/GPU pins that are associated with the ASIC. Further, an FPGA/GPU designer may implement any desired logic to command the ASIC, provide data to compute, and consume the results from the ASIC.

As discussed above, DNN algorithms provide a set of multiply/add operations on vector/matrix data. Additional operations, such as activation functions (e.g., ReLU) and scaling, are also commonly used. According to one embodiment, the FA accelerator provides support for such operations. In a further embodiment, the FA accelerator supports low precisions, including extremely low 1-bit and 2-bit operations for emerging binary and ternary DNNs, as well as efficient handling of sparse data by skipping unnecessary zero (or unimportant) values in sparse matrices. Additionally, the FA accelerator supports data compression to improve efficiency of data movements. FIG. **28F** illustrates another embodiment of an FA accelerator implemented in a DNN application. As shown in FIG. **28F**, the FA accelerator includes a systolic chain of processing elements (PEs) to perform computations **2833**, along with units at the input (**2832**) and output (**2836**) of the chain to manage efficient data movements between the ASIC and FPGA. (e.g., by appropriately sequencing and/or (de)compressing data).

FIG. **28G** illustrates one embodiment of a processing element, which includes a set of functional units (FUs) **2835** for computation, along with a scratchpad memory (SPAD) **2834** to feed the FUs **2835** with PE-local data. FIG. **28H** illustrates one embodiment of SPAD **2834**, which includes RAMs and support for compressed data such as weight sharing. FIG. **28I** illustrates one embodiment of a functional unit **2835** that includes variable precision multiply and add units, and sparse data scheduler that identifies unimportant values (e.g., zeros or certain range of values) and skips them.

In yet a further embodiment, accelerator **2708** is implemented to optimize batch-normalization layers to speedup overall network processing. Batch Normalization involves normalization of an input by subtracting a mean and dividing the variance calculated over a mini-batch during training (and inference). The batch normalization layer is now used in a majority of the start-of-the-art networks as it accelerates training by decreasing the internal covariate shift. Essentially, the batch normalization layer allows for larger learning rates, decreases sensitivity to initialization of weights, and in some cases aids in increase of accuracy. The advantages of using this layer however comes at a price in terms of computation and speed. Particularly, calculation of mean and variance for every pixel/input is computationally expen-

45

sive, making batch normalization the next compute intensive layer after convolution and fully connected layers.

As shown in FIG. 29, normalization operations occur only after convolution operations and the calculation of mean and the mean and variance calculation. According to one embodiment, accelerator 2708 provides dedicated hardware to calculate mean and variance operations, which enables acceleration of the batch normalization layer by enabling parallelization. FIG. 30 illustrates one embodiment in which a convolution operation 3002 is performed on input data, and the results are forwarded for parallel computation of mean and variance operations 3003 and normalization operations 3004. Once calculated, the results of mean and variance operations 3003 are forwarded to complete computation of normalization operations 3004.

FIG. 31 illustrates another embodiment in which accelerator 2708 features the fusion of the convolution operation and mean and variance calculation (3102). The result of those operations are subsequently forwarded for normalization operations 3004. This embodiment increases performance by decreasing memory access and running the required operations in fewer clock cycles.

The above-described batch-normalization embodiments accelerate training time in cases where mean and variance are calculated for a mini-batch during inference, as well as decrease inference. In other embodiments, normalization operations may also be diverted to accelerator 2708 so that the entire operation can be sped up. Since the operation happens on continuous memory location, hardware level parallelization has an efficient impact in improving layer performance.

FIG. 32 is a block diagram illustrating a computing system 3200 configured to implement one or more aspects of the embodiments described herein. The computing system 3200 includes a processing subsystem 3201 having one or more processor(s) 3212 and a system memory 3204 communicating via an interconnection path that may include a memory hub 3205. The memory hub 3205 may be a separate component within a chipset component or may be integrated within the one or more processor(s) 3202. The memory hub 3205 couples with an I/O subsystem 3211 via a communication link 3206. The I/O subsystem 3211 includes an I/O hub 3207 that can enable the computing system 3200 to receive input from one or more input device(s) 3208. Additionally, the I/O hub 3207 can enable a display controller, which may be included in the one or more processor(s) 3202, to provide outputs to one or more display device(s) 3210A. In one embodiment the one or more display device(s) 3210A coupled with the I/O hub 3207 can include a local, internal, or embedded display device.

In one embodiment the processing subsystem 3201 includes one or more parallel processor(s) 3212 coupled to memory hub 3205 via a bus or other communication link 3213. The communication link 3213 may be one of any number of standards based communication link technologies or protocols, such as, but not limited to PCI Express, or may be a vendor specific communications interface or communications fabric. In one embodiment the one or more parallel processor(s) 3212 form a computationally focused parallel or vector processing system that include a large number of processing cores and/or processing clusters, such as a many integrated core (MIC) processor. In one embodiment the one or more parallel processor(s) 3212 form a graphics processing subsystem that can output pixels to one of the one or more display device(s) 3210A coupled via the I/O Hub 3207. The one or more parallel processor(s) 3212 can also

46

include a display controller and display interface (not shown) to enable a direct connection to one or more display device(s) 3210B.

Within the I/O subsystem 3211, a system storage unit 3214 can connect to the I/O hub 3207 to provide a storage mechanism for the computing system 3200. An I/O switch 3216 can be used to provide an interface mechanism to enable connections between the I/O hub 3207 and other components, such as a network adapter 3218 and/or wireless network adapter 3219 that may be integrated into the platform, and various other devices that can be added via one or more add-in device(s) 3220. The network adapter 3218 can be an Ethernet adapter or another wired network adapter. The wireless network adapter 3219 can include one or more of a Wi-Fi, Bluetooth, near field communication (NFC), or other network device that includes one or more wireless radios.

The computing system 3200 can include other components not explicitly shown, including USB or other port connections, optical storage drives, video capture devices, and the like, may also be connected to the I/O hub 3207. Communication paths interconnecting the various components in FIG. 32 may be implemented using any suitable protocols, such as PCI (Peripheral Component Interconnect) based protocols (e.g., PCI-Express), or any other bus or point-to-point communication interfaces and/or protocol(s), such as the NV-Link high-speed interconnect, or interconnect protocols known in the art.

In one embodiment, the one or more parallel processor(s) 3212 incorporate circuitry optimized for graphics and video processing, including, for example, video output circuitry, and constitutes a graphics processing unit (GPU). In another embodiment, the one or more parallel processor(s) 3212 incorporate circuitry optimized for general purpose processing, while preserving the underlying computational architecture, described in greater detail herein. In yet another embodiment, components of the computing system 3200 may be integrated with one or more other system elements on a single integrated circuit. For example, the one or more parallel processor(s), 3212 memory hub 3205, processor(s) 3202, and I/O hub 3207 can be integrated into a system on chip (SoC) integrated circuit. Alternatively, the components of the computing system 3200 can be integrated into a single package to form a system in package (SIP) configuration. In one embodiment at least a portion of the components of the computing system 3200 can be integrated into a multi-chip module (MCM), which can be interconnected with other multi-chip modules into a modular computing system.

It will be appreciated that the computing system 3200 shown herein is illustrative and that variations and modifications are possible. The connection topology, including the number and arrangement of bridges, the number of processor(s) 3202, and the number of parallel processor(s) 3212, may be modified as desired. For instance, in some embodiments, system memory 3204 is connected to the processor(s) 3202 directly rather than through a bridge, while other devices communicate with system memory 3204 via the memory hub 3205 and the processor(s) 3202. In other alternative topologies, the parallel processor(s) 3212 are connected to the I/O hub 3207 or directly to one of the one or more processor(s) 3202, rather than to the memory hub 3205. In other embodiments, the I/O hub 3207 and memory hub 3205 may be integrated into a single chip. Some embodiments may include two or more sets of processor(s) 3202 attached via multiple sockets, which can couple with two or more instances of the parallel processor(s) 3212.

Some of the particular components shown herein are optional and may not be included in all implementations of the computing system **3200**. For example, any number of add-in cards or peripherals may be supported, or some components may be eliminated. Furthermore, some architectures may use different terminology for components similar to those illustrated in FIG. **32**. For example, the memory hub **3205** may be referred to as a Northbridge in some architectures, while the I/O hub **3207** may be referred to as a Southbridge.

FIG. **33A** illustrates a parallel processor **3300**, according to an embodiment. The various components of the parallel processor **3300** may be implemented using one or more integrated circuit devices, such as programmable processors, application specific integrated circuits (ASICs), or field programmable gate arrays (FPGA). The illustrated parallel processor **3300** is a variant of the one or more parallel processor(s) **3212** shown in FIG. **32**, according to an embodiment.

In one embodiment the parallel processor **3300** includes a parallel processing unit **3302**. The parallel processing unit includes an I/O unit **3304** that enables communication with other devices, including other instances of the parallel processing unit **3302**. The I/O unit **3304** may be directly connected to other devices. In one embodiment the I/O unit **3304** connects with other devices via the use of a hub or switch interface, such as memory hub **3205**. The connections between the memory hub **3205** and the I/O unit **3304** form a communication link **3213**. Within the parallel processing unit **3302**, the I/O unit **3304** connects with a host interface **3306** and a memory crossbar **3316**, where the host interface **3306** receives commands directed to performing processing operations and the memory crossbar **3316** receives commands directed to performing memory operations.

When the host interface **3306** receives a command buffer via the I/O unit **3304**, the host interface **3306** can direct work operations to perform those commands to a front end **3308**. In one embodiment the front end **3308** couples with a scheduler **3310**, which is configured to distribute commands or other work items to a processing cluster array **3312**. In one embodiment the scheduler **3310** ensures that the processing cluster array **3312** is properly configured and in a valid state before tasks are distributed to the processing clusters of the processing cluster array **3312**. In one embodiment the scheduler **3310** is implemented via firmware logic executing on a microcontroller. The microcontroller implemented scheduler **3310** is configurable to perform complex scheduling and work distribution operations at coarse and fine granularity, enabling rapid preemption and context switching of threads executing on the processing array **3312**. In one embodiment, the host software can prove workloads for scheduling on the processing array **3312** via one of multiple graphics processing doorbells. The workloads can then be automatically distributed across the processing array **3312** by the scheduler **3310** logic within the scheduler microcontroller.

The processing cluster array **3312** can include up to “N” processing clusters (e.g., cluster **3314A**, cluster **3314B**, through cluster **3314N**). Each cluster **3314A-3314N** of the processing cluster array **3312** can execute a large number of concurrent threads. The scheduler **3310** can allocate work to the clusters **3314A-3314N** of the processing cluster array **3312** using various scheduling and/or work distribution algorithms, which may vary depending on the workload arising for each type of program or computation. The scheduling can be handled dynamically by the scheduler

**3310**, or can be assisted in part by compiler logic during compilation of program logic configured for execution by the processing cluster array **3312**. In one embodiment, different clusters **3314A-3314N** of the processing cluster array **3312** can be allocated for processing different types of programs or for performing different types of computations.

The processing cluster array **3312** can be configured to perform various types of parallel processing operations. In one embodiment the processing cluster array **3312** is configured to perform general-purpose parallel compute operations. For example, the processing cluster array **3312** can include logic to execute processing tasks including filtering of video and/or audio data, performing modeling operations, including physics operations, and performing data transformations.

In one embodiment the processing cluster array **3312** is configured to perform parallel graphics processing operations. In embodiments in which the parallel processor **3300** is configured to perform graphics processing operations, the processing cluster array **3312** can include additional logic to support the execution of such graphics processing operations, including, but not limited to texture sampling logic to perform texture operations, as well as tessellation logic and other vertex processing logic. Additionally, the processing cluster array **3312** can be configured to execute graphics processing related shader programs such as, but not limited to vertex shaders, tessellation shaders, geometry shaders, and pixel shaders. The parallel processing unit **3302** can transfer data from system memory via the I/O unit **3304** for processing. During processing the transferred data can be stored to on-chip memory (e.g., parallel processor memory **3322**) during processing, then written back to system memory.

In one embodiment, when the parallel processing unit **3302** is used to perform graphics processing, the scheduler **3310** can be configured to divide the processing workload into approximately equal sized tasks, to better enable distribution of the graphics processing operations to multiple clusters **3314A-3314N** of the processing cluster array **3312**. In some embodiments, portions of the processing cluster array **3312** can be configured to perform different types of processing. For example, a first portion may be configured to perform vertex shading and topology generation, a second portion may be configured to perform tessellation and geometry shading, and a third portion may be configured to perform pixel shading or other screen space operations, to produce a rendered image for display. Intermediate data produced by one or more of the clusters **3314A-3314N** may be stored in buffers to allow the intermediate data to be transmitted between clusters **3314A-3314N** for further processing.

During operation, the processing cluster array **3312** can receive processing tasks to be executed via the scheduler **3310**, which receives commands defining processing tasks from front end **3308**. For graphics processing operations, processing tasks can include indices of data to be processed, e.g., surface (patch) data, primitive data, vertex data, and/or pixel data, as well as state parameters and commands defining how the data is to be processed (e.g., what program is to be executed). The scheduler **3310** may be configured to fetch the indices corresponding to the tasks or may receive the indices from the front end **3308**. The front end **3308** can be configured to ensure the processing cluster array **3312** is configured to a valid state before the workload specified by incoming command buffers (e.g., batch-buffers, push buffers, etc.) is initiated.

Each of the one or more instances of the parallel processing unit **3302** can couple with parallel processor memory **3322**. The parallel processor memory **3322** can be accessed via the memory crossbar **3316**, which can receive memory requests from the processing cluster array **3312** as well as the I/O unit **3304**. The memory crossbar **3316** can access the parallel processor memory **3322** via a memory interface **3318**. The memory interface **3318** can include multiple partition units (e.g., partition unit **3320A**, partition unit **3320B**, through partition unit **3320N**) that can each couple to a portion (e.g., memory unit) of parallel processor memory **3322**. In one implementation, the number of partition units **3320A-3320N** is configured to be equal to the number of memory units, such that a first partition unit **3320A** has a corresponding first memory unit **3324A**, a second partition unit **3320B** has a corresponding memory unit **3324B**, and an Nth partition unit **3320N** has a corresponding Nth memory unit **3324N**. In other embodiments, the number of partition units **3320A-3320N** may not be equal to the number of memory devices.

In various embodiments, the memory units **3324A-3324N** can include various types of memory devices, including dynamic random-access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory. In one embodiment, the memory units **3324A-3324N** may also include 3D stacked memory, including but not limited to high bandwidth memory (HBM). Persons skilled in the art will appreciate that the specific implementation of the memory units **3324A-3324N** can vary, and can be selected from one of various conventional designs. Render targets, such as frame buffers or texture maps may be stored across the memory units **3324A-3324N**, allowing partition units **3320A-3320N** to write portions of each render target in parallel to efficiently use the available bandwidth of parallel processor memory **3322**. In some embodiments, a local instance of the parallel processor memory **3322** may be excluded in favor of a unified memory design that utilizes system memory in conjunction with local cache memory.

In one embodiment, any one of the clusters **3314A-3314N** of the processing cluster array **3312** can process data that will be written to any of the memory units **3324A-3324N** within parallel processor memory **3322**. The memory crossbar **3316** can be configured to transfer the output of each cluster **3314A-3314N** to any partition unit **3320A-3320N** or to another cluster **3314A-3314N**, which can perform additional processing operations on the output. Each cluster **3314A-3314N** can communicate with the memory interface **3318** through the memory crossbar **3316** to read from or write to various external memory devices. In one embodiment the memory crossbar **3316** has a connection to the memory interface **3318** to communicate with the I/O unit **3304**, as well as a connection to a local instance of the parallel processor memory **3322**, enabling the processing units within the different processing clusters **3314A-3314N** to communicate with system memory or other memory that is not local to the parallel processing unit **3302**. In one embodiment the memory crossbar **3316** can use virtual channels to separate traffic streams between the clusters **3314A-3314N** and the partition units **3320A-3320N**.

While a single instance of the parallel processing unit **3302** is illustrated within the parallel processor **3300**, any number of instances of the parallel processing unit **3302** can be included. For example, multiple instances of the parallel processing unit **3302** can be provided on a single add-in card, or multiple add-in cards can be interconnected. The

different instances of the parallel processing unit **3302** can be configured to inter-operate even if the different instances have different numbers of processing cores, different amounts of local parallel processor memory, and/or other configuration differences. For example, in one embodiment some instances of the parallel processing unit **3302** can include higher precision floating point units relative to other instances. Systems incorporating one or more instances of the parallel processing unit **3302** or the parallel processor **3300** can be implemented in a variety of configurations and form factors, including but not limited to desktop, laptop, or handheld personal computers, servers, workstations, game consoles, and/or embedded systems.

FIG. **33B** is a block diagram of a partition unit **3320**, according to an embodiment. In one embodiment the partition unit **3320** is an instance of one of the partition units **3320A-3320N** of FIG. **33A**. As illustrated, the partition unit **3320** includes an L2 cache **221**, a frame buffer interface **3325**, and a ROP **3326** (raster operations unit). The L2 cache **221** is a read/write cache that is configured to perform load and store operations received from the memory crossbar **3316** and ROP **3326**. Read misses and urgent write-back requests are output by L2 cache **221** to frame buffer interface **3325** for processing. Updates can also be sent to the frame buffer via the frame buffer interface **3325** for processing. In one embodiment the frame buffer interface **3325** interfaces with one of the memory units in parallel processor memory, such as the memory units **3324A-3324N** of FIG. **33B** (e.g., within parallel processor memory **3322**).

In graphics applications, the ROP **3326** is a processing unit that performs raster operations such as stencil, z test, blending, and the like. The ROP **3326** then outputs processed graphics data that is stored in graphics memory. In some embodiments the ROP **3326** includes compression logic to compress depth or color data that is written to memory and decompress depth or color data that is read from memory. The compression logic can be lossless compression logic that makes use of one or more of multiple compression algorithms. The type of compression that is performed by the ROP **3326** can vary based on the statistical characteristics of the data to be compressed. For example, in one embodiment, delta color compression is performed on depth and color data on a per-tile basis.

In some embodiments, the ROP **3326** is included within each processing cluster (e.g., cluster **3314A-3314N** of FIG. **33A**) instead of within the partition unit **3320**. In such embodiment, read and write requests for pixel data are transmitted over the memory crossbar **3316** instead of pixel fragment data. The processed graphics data may be displayed on a display device, such as one of the one or more display device(s) **3210** of FIG. **32**, routed for further processing by the processor(s) **3302**, or routed for further processing by one of the processing entities within the parallel processor **3300** of FIG. **33A**.

FIG. **33C** is a block diagram of a processing cluster **3314** within a parallel processing unit, according to an embodiment. In one embodiment, the processing cluster is an instance of one of the processing clusters **3314A-3314N** of FIG. **33B**. The processing cluster **3314** can be configured to execute many threads in parallel, where the term "thread" refers to an instance of a particular program executing on a particular set of input data. In some embodiments, single-instruction, multiple-data (SIMD) instruction issue techniques are used to support parallel execution of a large number of threads without providing multiple independent instruction units. In other embodiments, single-instruction, multiple-thread (SIMT) techniques are used to support par-

## 51

allel execution of a large number of generally synchronized threads, using a common instruction unit configured to issue instructions to a set of processing engines within each one of the processing clusters. Unlike a SIMD execution regime, where all processing engines typically execute identical instructions, SIMT execution allows different threads to more readily follow divergent execution paths through a given thread program. Persons skilled in the art will understand that a SIMD processing regime represents a functional subset of a SIMT processing regime.

Operation of the processing cluster **3314** can be controlled via a pipeline manager **322** that distributes processing tasks to SIMT parallel processors. The pipeline manager **322** receives instructions from the scheduler **3310** of FIG. **33A** and manages execution of those instructions via a graphics multiprocessor **324** and/or a texture unit **326**. The illustrated graphics multiprocessor **324** is an exemplary instance of a SIMT parallel processor. However, various types of SIMT parallel processors of differing architectures may be included within the processing cluster **3314**. One or more instances of the graphics multiprocessor **3334** can be included within a processing cluster **3314**. The graphics multiprocessor **3334** can process data and a data crossbar **3340** can be used to distribute the processed data to one of multiple possible destinations, including other shader units. The pipeline manager **3332** can facilitate the distribution of processed data by specifying destinations for processed data to be distributed via the data crossbar **3340**.

Each graphics multiprocessor **3334** within the processing cluster **3314** can include an identical set of functional execution logic (e.g., arithmetic logic units, load-store units, etc.). The functional execution logic can be configured in a pipelined manner in which new instructions can be issued before previous instructions are complete. The functional execution logic supports a variety of operations including integer and floating-point arithmetic, comparison operations, Boolean operations, bit-shifting, and computation of various algebraic functions. In one embodiment the same functional-unit hardware can be leveraged to perform different operations and any combination of functional units may be present.

The instructions transmitted to the processing cluster **3314** constitutes a thread. A set of threads executing across the set of parallel processing engines is a thread group. A thread group executes the same program on different input data. Each thread within a thread group can be assigned to a different processing engine within a graphics multiprocessor **3334**. A thread group may include fewer threads than the number of processing engines within the graphics multiprocessor **3334**. When a thread group includes fewer threads than the number of processing engines, one or more of the processing engines may be idle during cycles in which that thread group is being processed. A thread group may also include more threads than the number of processing engines within the graphics multiprocessor **3334**. When the thread group includes more threads than the number of processing engines within the graphics multiprocessor **3334**, processing can be performed over consecutive clock cycles. In one embodiment multiple thread groups can be executed concurrently on a graphics multiprocessor **3334**. In one embodiment the graphics multiprocessor **3334** includes an internal cache memory to perform load and store operations. In one embodiment, the graphics multiprocessor **3334** can forego an internal cache and use a cache memory (e.g., L1 cache **3348**) within the processing cluster **3314**. Each graphics multiprocessor **3334** also has access to L2 caches within the partition units (e.g., partition units **3320A-3320N** of FIG.

## 52

**33B**) that are shared among all processing clusters **3314** and may be used to transfer data between threads. The graphics multiprocessor **3334** may also access off-chip global memory, which can include one or more of local parallel processor memory and/or system memory. Any memory external to the parallel processing unit **3302** may be used as global memory. Embodiments in which the processing cluster **3314** includes multiple instances of the graphics multiprocessor **3334** can share common instructions and data, which may be stored in the L1 cache **3348**.

Each processing cluster **3314** may include an MMU **3345** (memory management unit) that is configured to map virtual addresses into physical addresses. In other embodiments, one or more instances of the MMU **3345** may reside within the memory interface **3318** of FIG. **33A**. The MMU **3345** includes a set of page table entries (PTEs) used to map a virtual address to a physical address of a tile and optionally a cache line index. The MMU **3345** may include address translation lookaside buffers (TLB) or caches that may reside within the graphics multiprocessor **3334** or the L1 cache or processing cluster **3314**. The physical address is processed to distribute surface data access locality to allow efficient request interleaving among partition units. The cache line index may be used to determine whether a request for a cache line is a hit or miss. In graphics and computing applications, a processing cluster **3314** may be configured such that each graphics multiprocessor **3334** is coupled to a texture unit **3336** for performing texture mapping operations, e.g., determining texture sample positions, reading texture data, and filtering the texture data. Texture data is read from an internal texture L1 cache (not shown) or in some embodiments from the L1 cache within graphics multiprocessor **3334** and is fetched from an L2 cache, local parallel processor memory, or system memory, as needed. Each graphics multiprocessor **3334** outputs processed tasks to the data crossbar **3340** to provide the processed task to another processing cluster **3314** for further processing or to store the processed task in an L2 cache, local parallel processor memory, or system memory via the memory crossbar **3316**. A preROP **3342** (pre-raster operations unit) is configured to receive data from graphics multiprocessor **234**, direct data to ROP units, which may be located with partition units as described herein (e.g., partition units **3320A-3320N** of FIG. **33A**). The preROP **3342** unit can perform optimizations for color blending, organize pixel color data, and perform address translations.

It will be appreciated that the core architecture described herein is illustrative and that variations and modifications are possible. Any number of processing units, e.g., graphics multiprocessor **3334**, texture units **3336**, preROPs **3342**, etc., may be included within a processing cluster **3314**. Further, while only one processing cluster **3314** is shown, a parallel processing unit as described herein may include any number of instances of the processing cluster **3314**. In one embodiment, each processing cluster **3314** can be configured to operate independently of other processing clusters **3314** using separate and distinct processing units, L1 caches, etc.

FIG. **33D** shows a graphics multiprocessor **3334**, according to one embodiment. In such embodiment, the graphics multiprocessor **3334** couples with the pipeline manager **3332** of the processing cluster **3314**. The graphics multiprocessor **3334** has an execution pipeline including but not limited to an instruction cache **3352**, an instruction unit **3354**, an address mapping unit **3356**, a register file **3358**, one or more general purpose graphics processing unit (GPGPU) cores **3362**, and one or more load/store units **3366**. The

GPGPU cores **3362** and load/store units **3366** are coupled with cache memory **3372** and shared memory **3370** via a memory and cache interconnect **3368**.

In one embodiment, the instruction cache **3352** receives a stream of instructions to execute from the pipeline manager **3332**. The instructions are cached in the instruction cache **3352** and dispatched for execution by the instruction unit **3354**. The instruction unit **3354** can dispatch instructions as thread groups (e.g., warps), with each thread of the thread group assigned to a different execution unit within GPGPU core **3362**. An instruction can access any of a local, shared, or global address space by specifying an address within a unified address space. The address mapping unit **3356** can be used to translate addresses in the unified address space into a distinct memory address that can be accessed by the load/store units **3366**.

The register file **3358** provides a set of registers for the functional units of the graphics multiprocessor **3334**. The register file **3358** provides temporary storage for operands connected to the data paths of the functional units (e.g., GPGPU cores **3362**, load/store units **3366**) of the graphics multiprocessor **3334**. In one embodiment, the register file **3358** is divided between each of the functional units such that each functional unit is allocated a dedicated portion of the register file **3358**. In one embodiment, the register file **3358** is divided between the different warps being executed by the graphics multiprocessor **3334**.

The GPGPU cores **3362** can each include floating point units (FPUs) and/or integer arithmetic logic units (ALUs) that are used to execute instructions of the graphics multiprocessor **3334**. The GPGPU cores **3362** can be similar in architecture or can differ in architecture, according to embodiments. For example, in one embodiment, a first portion of the GPGPU cores **3362** include a single precision FPU and an integer ALU while a second portion of the GPGPU cores include a double precision FPU. In one embodiment, the FPUs can implement the IEEE 754-33008 standard for floating point arithmetic or enable variable precision floating point arithmetic. The graphics multiprocessor **3334** can additionally include one or more fixed function or special function units to perform specific functions such as copy rectangle or pixel blending operations. In one embodiment one or more of the GPGPU cores can also include fixed or special function logic.

In one embodiment, the GPGPU cores **3362** include SIMD logic capable of performing a single instruction on multiple sets of data. In one embodiment GPGPU cores **3362** can physically execute SIMD4, SIMD8, and SIMD16 instructions and logically execute SIMD1, SIMD2, and SIMD32 instructions. The SIMD instructions for the GPGPU cores can be generated at compile time by a shader compiler or automatically generated when executing programs written and compiled for single program multiple data (SPMD) or SIMT architectures. Multiple threads of a program configured for the SIMT execution model can be executed via a single SIMD instruction. For example and in one embodiment, eight SIMT threads that perform the same or similar operations can be executed in parallel via a single SIMD8 logic unit.

The memory and cache interconnect **3368** is an interconnect network that connects each of the functional units of the graphics multiprocessor **3334** to the register file **3358** and to the shared memory **3370**. In one embodiment, the memory and cache interconnect **3368** is a crossbar interconnect that allows the load/store unit **3366** to implement load and store operations between the shared memory **3370** and the register file **3358**. The register file **3358** can operate at the same

frequency as the GPGPU cores **3362**, thus data transfer between the GPGPU cores **3362** and the register file **3358** is very low latency. The shared memory **3370** can be used to enable communication between threads that execute on the functional units within the graphics multiprocessor **3334**. The cache memory **3372** can be used as a data cache for example, to cache texture data communicated between the functional units and the texture unit **3336**. The shared memory **3370** can also be used as a program managed cache. Threads executing on the GPGPU cores **3362** can programmatically store data within the shared memory in addition to the automatically cached data that is stored within the cache memory **3372**.

FIGS. **34A-34B** illustrate additional graphics multiprocessors, according to embodiments. The illustrated graphics multiprocessors **3425**, **3450** are variants of the graphics multiprocessor **3334** of FIG. **33C**. The illustrated graphics multiprocessors **3425**, **3450** can be configured as a streaming multiprocessor (SM) capable of simultaneous execution of a large number of execution threads.

FIG. **34A** shows a graphics multiprocessor **3425** according to an additional embodiment. The graphics multiprocessor **3425** includes multiple additional instances of execution resource units relative to the graphics multiprocessor **3334** of FIG. **33D**. For example, the graphics multiprocessor **3425** can include multiple instances of the instruction unit **3432A-3432B**, register file **3434A-3434B**, and texture unit(s) **3444A-3444B**. The graphics multiprocessor **3425** also includes multiple sets of graphics or compute execution units (e.g., GPGPU core **3436A-3436B**, GPGPU core **3437A-3437B**, GPGPU core **3438A-3438B**) and multiple sets of load/store units **3440A-3440B**. In one embodiment the execution resource units have a common instruction cache **3430**, texture and/or data cache memory **3442**, and shared memory **3446**.

The various components can communicate via an interconnect fabric **3427**. In one embodiment the interconnect fabric **3427** includes one or more crossbar switches to enable communication between the various components of the graphics multiprocessor **3425**. In one embodiment the interconnect fabric **3427** is a separate, high-speed network fabric layer upon which each component of the graphics multiprocessor **3425** is stacked. The components of the graphics multiprocessor **3425** communicate with remote components via the interconnect fabric **3427**. For example, the GPGPU cores **3436A-3436B**, **3437A-3437B**, and **34378A-3438B** can each communicate with shared memory **3446** via the interconnect fabric **3427**. The interconnect fabric **3427** can arbitrate communication within the graphics multiprocessor **3425** to ensure a fair bandwidth allocation between components.

FIG. **34B** shows a graphics multiprocessor **3450** according to an additional embodiment. The graphics processor includes multiple sets of execution resources **3456A-3456D**, where each set of execution resource includes multiple instruction units, register files, GPGPU cores, and load store units, as illustrated in FIG. **33D** and FIG. **34A**. The execution resources **3456A-3456D** can work in concert with texture unit(s) **3460A-3460D** for texture operations, while sharing an instruction cache **3454**, and shared memory **3462**. In one embodiment the execution resources **3456A-3456D** can share an instruction cache **3454** and shared memory **3462**, as well as multiple instances of a texture and/or data cache memory **3458A-3458B**. The various components can communicate via an interconnect fabric **3452** similar to the interconnect fabric **3427** of FIG. **34A**.

55

Persons skilled in the art will understand that the architecture described in FIGS. 32, 33A-33D, and 34A-34B are descriptive and not limiting as to the scope of the present embodiments. Thus, the techniques described herein may be implemented on any properly configured processing unit, including, without limitation, one or more mobile application processors, one or more desktop or server central processing units (CPUs) including multi-core CPUs, one or more parallel processing units, such as the parallel processing unit 3302 of FIG. 33, as well as one or more graphics processors or special purpose processing units, without departure from the scope of the embodiments described herein.

In some embodiments a parallel processor or GPGPU as described herein is communicatively coupled to host/processor cores to accelerate graphics operations, machine-learning operations, pattern analysis operations, and various general purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host processor/cores over a bus or other interconnect (e.g., a high speed interconnect such as PCIe or NVLink). In other embodiments, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

#### Techniques for GPU to Host Processor Interconnection

FIG. 35A illustrates an exemplary architecture in which a plurality of GPUs 3510-3513 are communicatively coupled to a plurality of multi-core processors 3505-3506 over high-speed links 3540-3543 (e.g., buses, point-to-point interconnects, etc.). In one embodiment, the high-speed links 3540-3543 support a communication throughput of 4 GB/s, 30 GB/s, 80 GB/s or higher, depending on the implementation. Various interconnect protocols may be used including, but not limited to, PCIe 4.0 or 5.0 and NVLink 2.0. However, the underlying principles of the invention are not limited to any particular communication protocol or throughput.

In addition, in one embodiment, two or more of the GPUs 3510-3513 are interconnected over high-speed links 3544-3545, which may be implemented using the same or different protocols/links than those used for high-speed links 3540-3543. Similarly, two or more of the multi-core processors 3505-3506 may be connected over high speed link 3533 which may be symmetric multi-processor (SMP) buses operating at 20 GB/s, 30 GB/s, 33320 GB/s or higher. Alternatively, all communication between the various system components shown in FIG. 35A may be accomplished using the same protocols/links (e.g., over a common interconnection fabric). As mentioned, however, the underlying principles of the invention are not limited to any particular type of interconnect technology.

In one embodiment, each multi-core processor 3505-3506 is communicatively coupled to a processor memory 3501-3502, via memory interconnects 3530-3531, respectively, and each GPU 3510-3513 is communicatively coupled to GPU memory 3520-3523 over GPU memory interconnects 3550-3553, respectively. The memory interconnects 3530-3531 and 3550-3553 may utilize the same or different memory access technologies. By way of example, and not limitation, the processor memories 3501-3502 and GPU memories 3520-3523 may be volatile memories such as

56

dynamic random access memories (DRAMs) (including stacked DRAMs), Graphics DDR SDRAM (GDDR) (e.g., GDDR5, GDDR6), or High Bandwidth Memory (HBM) and/or may be non-volatile memories such as 3D XPoint or Nano-Ram. In one embodiment, some portion of the memories may be volatile memory and another portion may be non-volatile memory (e.g., using a two-level memory (2LM) hierarchy).

As described below, although the various processors 3505-3506 and GPUs 3510-3513 may be physically coupled to a particular memory 3501-3502, 3520-3523, respectively, a unified memory architecture may be implemented in which the same virtual system address space (also referred to as the “effective address” space) is distributed among all of the various physical memories. For example, processor memories 3501-3502 may each comprise 64 GB of the system memory address space and GPU memories 3520-3523 may each comprise 32 GB of the system memory address space (resulting in a total of 256 GB addressable memory in this example).

FIG. 35B illustrates additional details for an interconnection between a multi-core processor 3507 and a graphics acceleration module 3546 in accordance with one embodiment. The graphics acceleration module 3546 may include one or more GPU chips integrated on a line card which is coupled to the processor 3507 via the high-speed link 3540. Alternatively, the graphics acceleration module 3546 may be integrated on the same package or chip as the processor 3507.

The illustrated processor 3507 includes a plurality of cores 3560A-3560D, each with a translation lookaside buffer 3561A-3561D and one or more caches 3562A-3562D. The cores may include various other components for executing instructions and processing data which are not illustrated to avoid obscuring the underlying principles of the invention (e.g., instruction fetch units, branch prediction units, decoders, execution units, reorder buffers, etc.). The caches 3562A-3562D may comprise level 1 (L1) and level 2 (L2) caches. In addition, one or more shared caches 3526 may be included in the caching hierarchy and shared by sets of the cores 3560A-3560D. For example, one embodiment of the processor 3507 includes 24 cores, each with its own L1 cache, twelve shared L2 caches, and twelve shared L3 caches. In this embodiment, one of the L2 and L3 caches are shared by two adjacent cores. The processor 3507 and the graphics accelerator integration module 3546 connect with system memory 3541, which may include processor memories 3501-3502.

Coherency is maintained for data and instructions stored in the various caches 3562A-3562D, 3556 and system memory 3541 via inter-core communication over a coherence bus 3564. For example, each cache may have cache coherency logic/circuitry associated therewith to communicate to over the coherence bus 3564 in response to detected reads or writes to particular cache lines. In one implementation, a cache snooping protocol is implemented over the coherence bus 3564 to snoop cache accesses. Cache snooping/coherency techniques are well understood by those of skill in the art and will not be described in detail here to avoid obscuring the underlying principles of the invention.

In one embodiment, a proxy circuit 3525 communicatively couples the graphics acceleration module 3546 to the coherence bus 3564, allowing the graphics acceleration module 3546 to participate in the cache coherence protocol as a peer of the cores. In particular, an interface 3535 provides connectivity to the proxy circuit 3525 over high-

57

speed link **3540** (e.g., a PCIe bus, NVLink, etc.) and an interface **3537** connects the graphics acceleration module **3546** to the link **3540**.

In one implementation, an accelerator integration circuit **3536** provides cache management, memory access, context management, and interrupt management services on behalf of a plurality of graphics processing engines **3531**, **3532**, N of the graphics acceleration module **3546**. The graphics processing engines **3531**, **3532**, N may each comprise a separate graphics processing unit (GPU). Alternatively, the graphics processing engines **3531**, **3532**, N may comprise different types of graphics processing engines within a GPU such as graphics execution units, media processing engines (e.g., video encoders/decoders), samplers, and blit engines. In other words, the graphics acceleration module may be a GPU with a plurality of graphics processing engines **3531**-**3532**, N or the graphics processing engines **3531**-**3532**, N may be individual GPUs integrated on a common package, line card, or chip.

In one embodiment, the accelerator integration circuit **3536** includes a memory management unit (MMU) **3539** for performing various memory management functions such as virtual-to-physical memory translations (also referred to as effective-to-real memory translations) and memory access protocols for accessing system memory **3541**. The MMU **3539** may also include a translation lookaside buffer (TLB) (not shown) for caching the virtual/effective to physical/real address translations. In one implementation, a cache **3538** stores commands and data for efficient access by the graphics processing engines **3531**-**3532**, N. In one embodiment, the data stored in cache **438** and graphics memories **3533**-**3534**, N is kept coherent with the core caches **3562A**-**3562D**, **3556** and system memory **3511**. As mentioned, this may be accomplished via proxy circuit **3525** which takes part in the cache coherency mechanism on behalf of cache **3538** and memories **3533**-**3534**, N (e.g., sending updates to the cache **3538** related to modifications/accesses of cache lines on processor caches **3562A**-**3562D**, **3556** and receiving updates from the cache **3538**).

A set of registers **3545** store context data for threads executed by the graphics processing engines **3531**-**3532**, N and a context management circuit **3548** manages the thread contexts. For example, the context management circuit **3548** may perform save and restore operations to save and restore contexts of the various threads during contexts switches (e.g., where a first thread is saved and a second thread is stored so that the second thread can be executed by a graphics processing engine). For example, on a context switch, the context management circuit **3548** may store current register values to a designated region in memory (e.g., identified by a context pointer). It may then restore the register values when returning to the context. In one embodiment, an interrupt management circuit **3547** receives and processes interrupts received from system devices.

In one implementation, virtual/effective addresses from a graphics processing engine **3531** are translated to real/physical addresses in system memory **3511** by the MMU **3539**. One embodiment of the accelerator integration circuit **3536** supports multiple (e.g., 4, 8, 16) graphics accelerator modules **3546** and/or other accelerator devices. The graphics accelerator module **3546** may be dedicated to a single application executed on the processor **3507** or may be shared between multiple applications. In one embodiment, a virtualized graphics execution environment is presented in which the resources of the graphics processing engines **3531**-**3532**, N are shared with multiple applications or virtual machines (VMs). The resources may be subdivided into "slices" which

58

are allocated to different VMs and/or applications based on the processing requirements and priorities associated with the VMs and/or applications.

Thus, the accelerator integration circuit acts as a bridge to the system for the graphics acceleration module **3546** and provides address translation and system memory cache services. In addition, the accelerator integration circuit **3536** may provide virtualization facilities for the host processor to manage virtualization of the graphics processing engines, interrupts, and memory management.

Because hardware resources of the graphics processing engines **3531**-**3532**, N are mapped explicitly to the real address space seen by the host processor **3507**, any host processor can address these resources directly using an effective address value. One function of the accelerator integration circuit **3536**, in one embodiment, is the physical separation of the graphics processing engines **3531**-**3532**, N so that they appear to the system as independent units.

As mentioned, in the illustrated embodiment, one or more graphics memories **3533**-**3534**, M are coupled to each of the graphics processing engines **3531**-**3532**, N, respectively. The graphics memories **3533**-**3534**, M store instructions and data being processed by each of the graphics processing engines **3531**-**3532**, N. The graphics memories **3533**-**3534**, M may be volatile memories such as DRAMs (including stacked DRAMs), GDDR memory (e.g., GDDR5, GDDR6), or HBM, and/or may be non-volatile memories such as 3D XPoint or Nano-Ram.

In one embodiment, to reduce data traffic over link **3540**, biasing techniques are used to ensure that the data stored in graphics memories **3533**-**3534**, M is data which will be used most frequently by the graphics processing engines **3531**-**3532**, N and preferably not used by the cores **3560A**-**3560D** (at least not frequently). Similarly, the biasing mechanism attempts to keep data needed by the cores (and preferably not the graphics processing engines **3531**-**3532**, N) within the caches **3562A**-**3562D**, **3556** of the cores and system memory **3511**.

FIG. **35C** illustrates another embodiment in which the accelerator integration circuit **3536** is integrated within the processor **3507**. In this embodiment, the graphics processing engines **3531**-**3532**, N communicate directly over the high-speed link **3540** to the accelerator integration circuit **3536** via interface **3537** and interface **3535** (which, again, may be utilize any form of bus or interface protocol). The accelerator integration circuit **3536** may perform the same operations as those described with respect to FIG. **35B**, but potentially at a higher throughput given its close proximity to the coherency bus **3562** and caches **3562A**-**3562D**, **3526**.

One embodiment supports different programming models including a dedicated-process programming model (no graphics acceleration module virtualization) and shared programming models (with virtualization). The latter may include programming models which are controlled by the accelerator integration circuit **3536** and programming models which are controlled by the graphics acceleration module **3546**. In one embodiment of the dedicated process model, graphics processing engines **3531**-**3532**, N are dedicated to a single application or process under a single operating system. The single application can funnel other application requests to the graphics engines **3531**-**3532**, N, providing virtualization within a VM/partition.

In the dedicated-process programming models, the graphics processing engines **3531**-**3532**, N, may be shared by multiple VM/application partitions. The shared models require a system hypervisor to virtualize the graphics processing engines **3531**-**3532**, N to allow access by each

59

operating system. For single-partition systems without a hypervisor, the graphics processing engines **3531-3532**, N are owned by the operating system. In both cases, the operating system can virtualize the graphics processing engines **3531-3532**, N to provide access to each process or application.

For the shared programming model, the graphics acceleration module **3546** or an individual graphics processing engine **3531-3532**, N selects a process element using a process handle. In one embodiment, process elements are stored in system memory **3511** and are addressable using the effective address to real address translation techniques described herein. The process handle may be an implementation-specific value provided to the host process when registering its context with the graphics processing engine **3531-3532**, N (that is, calling system software to add the process element to the process element linked list). The lower 16-bits of the process handle may be the offset of the process element within the process element linked list.

FIG. **35D** illustrates an exemplary accelerator integration slice **3590**. As used herein, a “slice” comprises a specified portion of the processing resources of the accelerator integration circuit **3536**. Application effective address space **3582** within system memory **3511** stores process elements **3583**. In one embodiment, the process elements **3583** are stored in response to GPU invocations **3581** from applications **3580** executed on the processor **3507**. A process element **3583** contains the process state for the corresponding application **3580**. A work descriptor (WD) **3584** contained in the process element **3583** can be a single job requested by an application or may contain a pointer to a queue of jobs. In the latter case, the WD **3584** is a pointer to the job request queue in the application’s address space **3582**.

The graphics acceleration module **3546** and/or the individual graphics processing engines **3531-3532**, N can be shared by all or a subset of the processes in the system. Embodiments of the invention include an infrastructure for setting up the process state and sending a WD **3584** to a graphics acceleration module **3546** to start a job in a virtualized environment.

In one implementation, the dedicated-process programming model is implementation-specific. In this model, a single process owns the graphics acceleration module **3546** or an individual graphics processing engine **3531**. Because the graphics acceleration module **3546** is owned by a single process, the hypervisor initializes the accelerator integration circuit **3536** for the owning partition and the operating system initializes the accelerator integration circuit **3536** for the owning process at the time when the graphics acceleration module **3546** is assigned.

In operation, a WD fetch unit **3591** in the accelerator integration slice **3590** fetches the next WD **3584** which includes an indication of the work to be done by one of the graphics processing engines of the graphics acceleration module **3546**. Data from the WD **3584** may be stored in registers **3545** and used by the MMU **3539**, interrupt management circuit **3547** and/or context management circuit **3546** as illustrated. For example, one embodiment of the MMU **3539** includes segment/page walk circuitry for accessing segment/page tables **3586** within the OS virtual address space **3585**. The interrupt management circuit **3547** may process interrupt events **3592** received from the graphics acceleration module **3546**. When performing graphics operations, an effective address **3593** generated by a graphics processing engine **3531-3532**, N is translated to a real address by the MMU **3539**.

60

In one embodiment, the same set of registers **3545** are duplicated for each graphics processing engine **3531-3532**, N and/or graphics acceleration module **3546** and may be initialized by the hypervisor or operating system. Each of these duplicated registers may be included in an accelerator integration slice **3590**. Exemplary registers that may be initialized by the hypervisor are shown in Table 1.

TABLE 1

## Hypervisor Initialized Registers

- |   |                                                                     |
|---|---------------------------------------------------------------------|
| 1 | Slice Control Register                                              |
| 2 | Real Address (RA) Scheduled Processes Area Pointer                  |
| 3 | Authority Mask Override Register                                    |
| 4 | Interrupt Vector Table Entry Offset                                 |
| 5 | Interrupt Vector Table Entry Limit                                  |
| 6 | State Register                                                      |
| 7 | Logical Partition ID                                                |
| 8 | Real address (RA) Hypervisor Accelerator Utilization Record Pointer |
| 9 | Storage Description Register                                        |

Exemplary registers that may be initialized by the operating system are shown in Table 2.

TABLE 2

## Operating System Initialized Registers

- |   |                                                             |
|---|-------------------------------------------------------------|
| 1 | Process and Thread Identification                           |
| 2 | Effective Address (EA) Context Save/Restore Pointer         |
| 3 | Virtual Address (VA) Accelerator Utilization Record Pointer |
| 4 | Virtual Address (VA) Storage Segment Table Pointer          |
| 5 | Authority Mask                                              |
| 6 | Work descriptor                                             |

In one embodiment, each WD **3584** is specific to a particular graphics acceleration module **446** and/or graphics processing engine **3531-3532**, N. It contains all the information a graphics processing engine **3531-3532**, N requires to do its work or it can be a pointer to a memory location where the application has set up a command queue of work to be completed.

FIG. **35E** illustrates additional details for one embodiment of a shared model. This embodiment includes a hypervisor real address space **3598** in which a process element list **3599** is stored. The hypervisor real address space **3598** is accessible via a hypervisor **3596** which virtualizes the graphics acceleration module engines for the operating system **3595**.

The shared programming models allow for all or a subset of processes from all or a subset of partitions in the system to use a graphics acceleration module **3546**. There are two programming models where the graphics acceleration module **3546** is shared by multiple processes and partitions: time-sliced shared and graphics directed shared.

In this model, the system hypervisor **3596** owns the graphics acceleration module **3546** and makes its function available to all operating systems **3595**. For a graphics acceleration module **3546** to support virtualization by the system hypervisor **3596**, the graphics acceleration module **3546** may adhere to the following requirements: 1) An application’s job request must be autonomous (that is, the state does not need to be maintained between jobs), or the graphics acceleration module **3546** must provide a context save and restore mechanism. 2) An application’s job request is guaranteed by the graphics acceleration module **3546** to complete in a specified amount of time, including any translation faults, or the graphics acceleration module **3546** provides the ability to preempt the processing of the job. 3)

## 61

The graphics acceleration module **3546** must be guaranteed fairness between processes when operating in the directed shared programming model.

In one embodiment, for the shared model, the application **3580** is required to make an operating system **3595** system call with a graphics acceleration module **3546** type, a work descriptor (WD), an authority mask register (AMR) value, and a context save/restore area pointer (CSRP). The graphics acceleration module **3546** type describes the targeted acceleration function for the system call. The graphics acceleration module **3546** type may be a system-specific value. The WD is formatted specifically for the graphics acceleration module **3546** and can be in the form of a graphics acceleration module **3546** command, an effective address pointer to a user-defined structure, an effective address pointer to a queue of commands, or any other data structure to describe the work to be done by the graphics acceleration module **3546**. In one embodiment, the AMR value is the AMR state to use for the current process. The value passed to the operating system is similar to an application setting the AMR. If the accelerator integration circuit **3536** and graphics acceleration module **3546** implementations do not support a User Authority Mask Override Register (UAMOR), the operating system may apply the current UAMOR value to the AMR value before passing the AMR in the hypervisor call. The hypervisor **3596** may optionally apply the current Authority Mask Override Register (AMOR) value before placing the AMR into the process element **3583**. In one embodiment, the CSRP is one of the registers **3545** containing the effective address of an area in the application's address space **3582** for the graphics acceleration module **3546** to save and restore the context state. This pointer is optional if no state is required to be saved between jobs or when a job is preempted. The context save/restore area may be pinned system memory.

Upon receiving the system call, the operating system **3595** may verify that the application **3580** has registered and been given the authority to use the graphics acceleration module **3546**. The operating system **3595** then calls the hypervisor **3596** with the information shown in Table 3.

TABLE 3

| OS to Hypervisor Call Parameters |                                                                      |
|----------------------------------|----------------------------------------------------------------------|
| 1                                | A work descriptor (WD)                                               |
| 2                                | An Authority Mask Register (AMR) value (potentially masked).         |
| 3                                | An effective address (EA) Context Save/Restore Area Pointer (CSRP)   |
| 4                                | A process ID (PID) and optional thread ID (TID)                      |
| 5                                | A virtual address (VA) accelerator utilization record pointer (AURP) |
| 6                                | The virtual address of the storage segment table pointer (SSTP)      |
| 7                                | A logical interrupt service number (LISN)                            |

Upon receiving the hypervisor call, the hypervisor **3596** verifies that the operating system **3595** has registered and been given the authority to use the graphics acceleration module **3546**. The hypervisor **3596** then puts the process element **3583** into the process element linked list for the corresponding graphics acceleration module **3546** type. The process element may include the information shown in Table 4.

TABLE 4

| Process Element Information |                                                              |
|-----------------------------|--------------------------------------------------------------|
| 1                           | A work descriptor (WD)                                       |
| 2                           | An Authority Mask Register (AMR) value (potentially masked). |

## 62

TABLE 4-continued

| Process Element Information |                                                                       |
|-----------------------------|-----------------------------------------------------------------------|
| 3                           | An effective address (EA) Context Save/Restore Area Pointer (CSRP)    |
| 4                           | A process ID (PID) and optional thread ID (TID)                       |
| 5                           | A virtual address (VA) accelerator utilization record pointer (AURP)  |
| 6                           | The virtual address of the storage segment table pointer (SSTP)       |
| 7                           | A logical interrupt service number (LISN)                             |
| 8                           | Interrupt vector table, derived from the hypervisor call parameters.  |
| 9                           | A state register (SR) value                                           |
| 10                          | A logical partition ID (LPID)                                         |
| 11                          | A real address (RA) hypervisor accelerator utilization record pointer |
| 12                          | The Storage Descriptor Register (SDR)                                 |

In one embodiment, the hypervisor initializes a plurality of accelerator integration slice **3590** registers **3545**.

As illustrated in FIG. **35F**, one embodiment of the invention employs a unified memory addressable via a common virtual memory address space used to access the physical processor memories **3501-3502** and GPU memories **3520-3523**. In this implementation, operations executed on the GPUs **3510-3513** utilize the same virtual/effective memory address space to access the processors memories **3501-3502** and vice versa, thereby simplifying programmability. In one embodiment, a first portion of the virtual/effective address space is allocated to the processor memory **3501**, a second portion to the second processor memory **3502**, a third portion to the GPU memory **3520**, and so on. The entire virtual/effective memory space (sometimes referred to as the effective address space) is thereby distributed across each of the processor memories **3501-3502** and GPU memories **3520-3523**, allowing any processor or GPU to access any physical memory with a virtual address mapped to that memory.

In one embodiment, bias/coherence management circuitry **3594A-3594E** within one or more of the MMUs **3539A-3539E** ensures cache coherence between the caches of the host processors (e.g., **3505**) and the GPUs **3510-3513** and implements biasing techniques indicating the physical memories in which certain types of data should be stored. While multiple instances of bias/coherence management circuitry **3594A-3594E** are illustrated in FIG. **35F**, the bias/coherence circuitry may be implemented within the MMU of one or more host processors **3505** and/or within the accelerator integration circuit **3536**.

One embodiment allows GPU-attached memory **3520-3523** to be mapped as part of system memory, and accessed using shared virtual memory (SVM) technology, but without suffering the typical performance drawbacks associated with full system cache coherence. The ability to GPU-attached memory **3520-3523** to be accessed as system memory without onerous cache coherence overhead provides a beneficial operating environment for GPU offload. This arrangement allows the host processor **3505** software to setup operands and access computation results, without the overhead of traditional I/O DMA data copies. Such traditional copies involve driver calls, interrupts and memory mapped I/O (MMIO) accesses that are all inefficient relative to simple memory accesses. At the same time, the ability to access GPU attached memory **3520-3523** without cache coherence overheads can be critical to the execution time of an offloaded computation. In cases with substantial streaming write memory traffic, for example, cache coherence overhead can significantly reduce the effective write bandwidth seen by a GPU **3510-3513**. The efficiency of operand setup, the efficiency of results access, and the efficiency of GPU computation all play a role in determining the effectiveness of GPU offload.

In one implementation, the selection of between GPU bias and host processor bias is driven by a bias tracker data structure. A bias table may be used, for example, which may be a page-granular structure (i.e., controlled at the granularity of a memory page) that includes 1 or 2 bits per GPU-attached memory page. The bias table may be implemented in a stolen memory range of one or more GPU-attached memories **3520-3523**, with or without a bias cache in the GPU **3510-3513** (e.g., to cache frequently/recently used entries of the bias table). Alternatively, the entire bias table may be maintained within the GPU.

In one implementation, the bias table entry associated with each access to the GPU-attached memory **3520-3523** is accessed prior the actual access to the GPU memory, causing the following operations. First, local requests from the GPU **3510-3513** that find their page in GPU bias are forwarded directly to a corresponding GPU memory **3520-3523**. Local requests from the GPU that find their page in host bias are forwarded to the processor **3505** (e.g., over a high-speed link as discussed above). In one embodiment, requests from the processor **3505** that find the requested page in host processor bias complete the request like a normal memory read. Alternatively, requests directed to a GPU-biased page may be forwarded to the GPU **3510-3513**. The GPU may then transition the page to a host processor bias if it is not currently using the page. The bias state of a page can be changed either by a software-based mechanism, a hardware-assisted software-based mechanism, or, for a limited set of cases, a purely hardware-based mechanism.

One mechanism for changing the bias state employs an API call (e.g. OpenCL), which, in turn, calls the GPU's device driver which, in turn, sends a message (or enqueues a command descriptor) to the GPU directing it to change the bias state and, for some transitions, perform a cache flushing operation in the host. The cache flushing operation is required for a transition from host processor **3505** bias to GPU bias, but is not required for the opposite transition.

In one embodiment, cache coherency is maintained by temporarily rendering GPU-biased pages uncacheable by the host processor **3505**. To access these pages, the processor **3505** may request access from the GPU **3510** which may or may not grant access right away, depending on the implementation. Thus, to reduce communication between the processor **3505** and GPU **3510** it is beneficial to ensure that GPU-biased pages are those which are required by the GPU but not the host processor **3505** and vice versa.

FIG. **35G** illustrates a multi-GPU computing system, according to an embodiment. The multi-GPU computing system can include a processor **3503** coupled to multiple GPUs **3514A-3514D** via a host interface switch **3504**. The host interface switch **3504**, in one embodiment, is a PCI express switch device that couples the processor **3503** to a PCI express bus over which the processor **3503** can communicate with the set of GPUs **3514A-3514D**. The GPUs **3514A-414D** can interconnect via a set of high-speed point to point GPU to GPU links **3516**. The high-speed GPU to GPU links can connect to each of the GPUs **3514A-3514D** via a dedicated GPU link. The P2P GPU links **3516** enable direct communication between each of the GPUs **3514A-3514D** without requiring communication over the host interface bus to which the processor **3503** is connected. With GPU-to-GPU traffic directed to the P2P GPU links, the host interface bus remains available for system memory access or to communicate with other instances of the multi-GPU computing system **3500**, for example, via one or more network devices. While in the illustrated embodiment the GPUs **3514A-3514D** connect to the processor **3503** via the

host interface switch **3504**, in one embodiment the processor **3503** includes direct support for the P2P GPU links **3516** and can connect directly to the GPUs **3514A-3514D**.

#### Graphics Processing Pipeline

FIG. **36** illustrates a graphics processing pipeline **3600**, according to an embodiment. In one embodiment, a graphics processor can implement the illustrated graphics processing pipeline **3600**. The graphics processor can be included within the parallel processing subsystems as described herein, such as the parallel processor **3300** of FIG. **33A**, which, in one embodiment, is a variant of the parallel processor(s) **3212** of FIG. **32**. The various parallel processing systems can implement the graphics processing pipeline **3600** via one or more instances of the parallel processing unit (e.g., parallel processing unit **3302** of FIG. **33A**) as described herein. For example, a shader unit (e.g., graphics multiprocessor **234** of FIG. **34**) may be configured to perform the functions of one or more of a vertex processing unit **3604**, a tessellation control processing unit **3608**, a tessellation evaluation processing unit **3612**, a geometry processing unit **3616**, and a fragment/pixel processing unit **3624**. The functions of data assembler **3602**, primitive assemblers **3606**, **3614**, **3618**, tessellation unit **3610**, rasterizer **3622**, and raster operations unit **3626** may also be performed by other processing engines within a processing cluster (e.g., processing cluster **3314** of FIG. **34**) and a corresponding partition unit (e.g., partition unit **3320A-3320N** of FIG. **33A**). The graphics processing pipeline **3600** may also be implemented using dedicated processing units for one or more functions. In one embodiment, one or more portions of the graphics processing pipeline **3600** can be performed by parallel processing logic within a general-purpose processor (e.g., CPU). In one embodiment, one or more portions of the graphics processing pipeline **3600** can access on-chip memory (e.g., parallel processor memory **3322** as in FIG. **33A**) via a memory interface **3628**, which may be an instance of the memory interface **3318** of FIG. **33A**.

In one embodiment, the data assembler **3602** is a processing unit that collects vertex data for surfaces and primitives. The data assembler **3602** then outputs the vertex data, including the vertex attributes, to the vertex processing unit **3604**. The vertex processing unit **3604** is a programmable execution unit that executes vertex shader programs, lighting and transforming vertex data as specified by the vertex shader programs. The vertex processing unit **3604** reads data that is stored in cache, local or system memory for use in processing the vertex data and may be programmed to transform the vertex data from an object-based coordinate representation to a world space coordinate space or a normalized device coordinate space.

A first instance of a primitive assembler **3606** receives vertex attributes from the vertex processing unit **3604**. The primitive assembler **3606** reads stored vertex attributes as needed and constructs graphics primitives for processing by tessellation control processing unit **3608**. The graphics primitives include triangles, line segments, points, patches, and so forth, as supported by various graphics processing application programming interfaces (APIs).

The tessellation control processing unit **3608** treats the input vertices as control points for a geometric patch. The control points are transformed from an input representation from the patch (e.g., the patch's bases) to a representation that is suitable for use in surface evaluation by the tessellation evaluation processing unit **3612**. The tessellation control processing unit **3608** can also compute tessellation factors for edges of geometric patches. A tessellation factor applies to a single edge and quantifies a view-dependent

65

level of detail associated with the edge. A tessellation unit **3610** is configured to receive the tessellation factors for edges of a patch and to tessellate the patch into multiple geometric primitives such as line, triangle, or quadrilateral primitives, which are transmitted to a tessellation evaluation processing unit **3612**. The tessellation evaluation processing unit **3612** operates on parameterized coordinates of the subdivided patch to generate a surface representation and vertex attributes for each vertex associated with the geometric primitives.

A second instance of a primitive assembler **3614** receives vertex attributes from the tessellation evaluation processing unit **3612**, reading stored vertex attributes as needed, and constructs graphics primitives for processing by the geometry processing unit **3616**. The geometry processing unit **3616** is a programmable execution unit that executes geometry shader programs to transform graphics primitives received from primitive assembler **3614** as specified by the geometry shader programs. In one embodiment the geometry processing unit **3616** is programmed to subdivide the graphics primitives into one or more new graphics primitives and calculate parameters used to rasterize the new graphics primitives.

In some embodiments the geometry processing unit **3616** can add or delete elements in the geometry stream. The geometry processing unit **3616** outputs the parameters and vertices specifying new graphics primitives to primitive assembler **3618**. The primitive assembler **3618** receives the parameters and vertices from the geometry processing unit **3616** and constructs graphics primitives for processing by a viewport scale, cull, and clip unit **3620**. The geometry processing unit **3616** reads data that is stored in parallel processor memory or system memory for use in processing the geometry data. The viewport scale, cull, and clip unit **3620** performs clipping, culling, and viewport scaling and outputs processed graphics primitives to a rasterizer **3622**. The rasterizer **3622** can perform depth culling and other depth-based optimizations. The rasterizer **3622** also performs scan conversion on the new graphics primitives to generate fragments and output those fragments and associated coverage data to the fragment/pixel processing unit **3624**. The fragment/pixel processing unit **3624** is a programmable execution unit that is configured to execute fragment shader programs or pixel shader programs. The fragment/pixel processing unit **3624** transforming fragments or pixels received from rasterizer **3622**, as specified by the fragment or pixel shader programs. For example, the fragment/pixel processing unit **3624** may be programmed to perform operations included but not limited to texture mapping, shading, blending, texture correction and perspective correction to produce shaded fragments or pixels that are output to a raster operations unit **3626**. The fragment/pixel processing unit **3624** can read data that is stored in either the parallel processor memory or the system memory for use when processing the fragment data. Fragment or pixel shader programs may be configured to shade at sample, pixel, tile, or other granularities depending on the sampling rate configured for the processing units.

The raster operations unit **3626** is a processing unit that performs raster operations including, but not limited to stencil, z test, blending, and the like, and outputs pixel data as processed graphics data to be stored in graphics memory (e.g., parallel processor memory **3322** as in FIG. 33A, and/or system memory **3204** as in FIG. 32, to be displayed on the one or more display device(s) **3210** or for further processing by one of the one or more processor(s) **3202** or parallel processor(s) **3212**. In some embodiments the raster opera-

66

tions unit **3626** is configured to compress z or color data that is written to memory and decompress z or color data that is read from memory.

The foregoing description and drawings are to be regarded in an illustrative rather than a restrictive sense. Persons skilled in the art will understand that various modifications and changes may be made to the embodiments described herein without departing from the broader spirit and scope of the invention as set forth in the appended claims.

Some embodiments pertain to Example 1 that includes an apparatus to facilitate acceleration of machine learning operations, comprising at least one processor to perform operations to implement a neural network and accelerator logic communicatively coupled to the processor to perform compute operations for the neural network.

Example 2 includes the subject matter of Example 1, wherein the at least one processor comprises a graphics processing unit (GPU) and the accelerator logic comprises a central processing unit (CPU).

Example 3 includes the subject matter of Examples 1 and 2, wherein the neural network comprises a Differentiable Neural Computer (DNC), comprising an external memory coupled to the CPU and the GPU to store knowledge data for the neural network.

Example 4 includes the subject matter of Examples 1-3, wherein the CPU performs transformation compute operations for the neural network.

Example 5 includes the subject matter of Examples 1-4, wherein the CPU performs a zero copy operation to facilitate a transfer of data between the CPU and the GPU.

Example 6 includes the subject matter of Examples 1-5, wherein the accelerator logic comprises a scheduler to analyze CPU and GPU resources and a compute graph of the neural network, assign nodes of the compute graph to the CPU and GPU resources and schedule the compute graph for processing at the CPU and GPU resources.

Example 7 includes the subject matter of Examples 1-6, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

Example 8 includes the subject matter of Examples 1-7, wherein the computation cost is determined based on individually processing the operators at the CPU and GPU.

Example 9 includes the subject matter of Examples 1-8, wherein the computation cost is determined based on simultaneously processing the operators at the CPU and GPU.

Example 10 includes the subject matter of Examples 1-9, wherein assigning the nodes of the compute graph to the CPU and GPU resources comprises determining a shortest path for each of the operators based on the computation cost.

Example 11 includes the subject matter of Examples 1-10, wherein the accelerator logic comprises an accelerator integrated circuit (IC) to perform data layout transformations.

Example 12 includes the subject matter of Examples 1-11, wherein the data layout transformations comprise computing weight gradient computations during neural network layer training.

Example 13 includes the subject matter of Examples 1-12, wherein the accelerator IC performs the weight gradient computations during back-propagation computation performed at the at least one processor.

Example 14 includes the subject matter of Examples 1-13, wherein the accelerator IC comprises a memory to store input matrix data received from the at least one processor during forward propagation computations during neural network layer training.

67

Example 15 includes the subject matter of Examples 1-14, wherein the accelerator IC performs a transpose computation on the input matrix data upon receiving the input matrix data in advance of performing the weight gradient computations.

Example 16 includes the subject matter of Examples 1-15, wherein the integrated circuit IC comprises a FPGA-ASIC (FA) integrated circuit architecture.

Some embodiments pertain to Example 17 that includes a method to facilitate acceleration of machine learning operations, analyzing central processing unit (CPU) and graphics processing unit (GPU) resources and a compute graph of a neural network, assigning nodes of the compute graph to the CPU and GPU resources and scheduling the compute graph for processing at the CPU and GPU resources.

Example 18 includes the subject matter of Example 17, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

Example 19 includes the subject matter of Examples 17 and 18, wherein the computation cost is determined based on individually processing the operators at the CPU and GPU.

Example 20 includes the subject matter of Examples 17-19, wherein the computation cost is determined based on simultaneously processing the operators at the CPU and GPU.

Example 21 includes the subject matter of Examples 17-20, wherein assigning the nodes of the compute graph to the CPU and GPU resources comprises determining a shortest path for each of the operators based on the computation cost.

The foregoing description and drawings are to be regarded in an illustrative rather than a restrictive sense. Persons skilled in the art will understand that various modifications and changes may be made to the embodiments described herein without departing from the broader spirit and scope of the invention as set forth in the appended claims.

What is claimed is:

1. An apparatus comprising:

a graphics processing unit (GPU) to perform operations to implement a neural network; and

accelerator circuitry to perform compute operations for the neural network using a Differentiable Neural Computer (DNC) neural network architecture implemented by the accelerator circuitry, wherein the accelerator circuitry to:

implement a fast interconnect between the GPU and a central processing unit (CPU) to accelerate deep learning tasks of the neural network;

analyze hardware resources of the CPU and the GPU and analyze a compute graph of the neural network;

assign, based on analysis of the hardware resources and of the compute graph, nodes of the compute graph to the CPU and GPU resources; and

schedule the compute graph for processing at the hardware resources of the CPU and of the GPU;

wherein the DNC neural network architecture comprises an external memory coupled to the CPU, the external memory comprising an auto associative memory to store knowledge data for the neural network.

2. The apparatus of claim 1, wherein the CPU performs transformation compute operations for the neural network.

3. The apparatus of claim 1, wherein the CPU performs a zero copy operation to facilitate a transfer of data between the CPU and the GPU.

68

4. The apparatus of claim 1, wherein CPU comprises the accelerator circuitry, and wherein the accelerator circuitry comprises a scheduler to analyze the hardware resources, assign the nodes, and schedule the compute graph for processing at the CPU and GPU resources.

5. The apparatus of claim 4, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

6. The apparatus of claim 5, wherein assigning the nodes of the compute graph to the CPU and GPU resources comprises determining a shortest path for each of the operators based on the computation cost.

7. The apparatus of claim 5, wherein the computation cost is determined based on at least one of individually processing the operators at the CPU and GPU or simultaneously processing the operators at the CPU and GPU.

8. The apparatus of claim 1, wherein the accelerator circuitry comprises an accelerator integrated circuit (IC) to perform data layout transformations, and wherein the data layout transformations comprise computing weight gradient computations during neural network layer training.

9. The apparatus of claim 8, wherein the accelerator IC performs the weight gradient computations during back-propagation computation.

10. The apparatus of claim 8, wherein the accelerator IC comprises a memory to store input matrix data received from at least one of the CPU or GPU during forward propagation computations during neural network layer training.

11. The apparatus of claim 10, wherein the accelerator IC performs a transpose computation on the input matrix data upon receiving the input matrix data in advance of performing the weight gradient computations.

12. The apparatus of claim 8, wherein the integrated circuit (IC) comprises a FPGA-ASIC (FA) IC architecture.

13. The apparatus of claim 8, wherein the IC comprises circuitry to calculate mean and variance operations to accelerate a batch normalization layer and wherein the calculation of the mean and variance operations is performed in parallel with normalization operations.

14. The apparatus of claim 13, wherein the calculation of the mean and variance operations is fused with a convolution operation.

15. A method comprising:

implementing, by accelerator circuitry, a fast interconnect between a graphics processing unit (GPU) and a central processing unit (CPU) to accelerate deep learning tasks of a neural network;

performing, by the accelerator circuitry, compute operations for the neural network using a Differentiable Neural Computer (DNC) neural network architecture implemented by the accelerator circuitry, the DNC neural network architecture comprises an external memory comprising an auto associative memory to store knowledge data for the neural network;

analyzing, by the accelerator circuitry, CPU and GPU resources and analyzing a compute graph of the neural network;

assigning, by the accelerator circuitry based on analysis of the CPU and GPU resources and of the compute graph, nodes of the compute graph to the CPU and GPU resources; and

scheduling, by the accelerator circuitry, the compute graph for processing at the CPU and GPU resources.

69

**16.** The method of claim **15**, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

**17.** The method of claim **16**, wherein the computation cost is determined based on individually processing the operators at the CPU and GPU.

**18.** The method of claim **16**, wherein the computation cost is determined based on simultaneously processing the operators at the CPU and GPU.

**19.** The method of claim **16**, wherein assigning the nodes of the compute graph to the CPU and GPU resources comprises determining a shortest path for each of the operators based on the computation cost.

**20.** A system comprising:

a memory;

a central processing unit (CPU);

a graphics processing unit (GPU) to perform operations to implement a neural network; and

accelerator circuitry to perform compute operations for the neural network using a Differentiable Neural Computer (DNC) neural network architecture implemented by the accelerator circuitry, wherein the accelerator circuitry to:

implement a fast interconnect between the GPU and the CPU to accelerate deep learning tasks of the neural network;

analyze hardware resources of the CPU and the GPU and analyze a compute graph of the neural network;

assign, based on analysis of the hardware resources and of the compute graph, nodes of the compute graph to the CPU and GPU resources; and

schedule the compute graph for processing at the hardware resources of the CPU and of the GPU;

wherein the DNC neural network architecture utilizes the memory as an external memory comprising an auto associative memory to store knowledge data for the neural network.

**21.** The system of claim **20**, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

70

**22.** The system of claim **21**, wherein the computation cost is determined based on at least one of individually processing the operators at the CPU and GPU or simultaneously processing the operators at the CPU and GPU.

**23.** A non-transitory computer-readable storage medium having stored thereon executable computer program instructions that, when executed by one or more processors, cause the one or more processors to perform operations comprising:

implementing, by accelerator circuitry of the one or more processors, a fast interconnect between a graphics processing unit (GPU) and a central processing unit (CPU) to accelerate deep learning tasks of a neural network;

performing, by the accelerator circuitry, compute operations for the neural network using a Differentiable Neural Computer (DNC) neural network architecture implemented by the accelerator circuitry, the DNC neural network architecture comprises an external memory comprising an auto associative memory to store knowledge data for the neural network;

analyzing, by the accelerator circuitry, CPU and GPU resources and analyzing a compute graph of the neural network;

assigning, by the accelerator circuitry based on analysis of the CPU and GPU resources and of the compute graph, nodes of the compute graph to the CPU and GPU resources; and

scheduling, by the accelerator circuitry, the compute graph for processing at the CPU and GPU resources.

**24.** The non-transitory computer-readable storage medium of claim **23**, wherein analyzing the CPU and GPU resources and the compute graph comprises determining a computation cost of operators to be performed at CPU and GPU.

**25.** The non-transitory computer-readable storage medium of claim **24**, wherein the computation cost is determined based on at least one of individually processing the operators at the CPU and GPU or simultaneously processing the operators at the CPU and GPU.

\* \* \* \* \*